

Expert/Novice differences in Diagnostic Medical Cognition - A Review of the Literature

**L. Cuthbert ¹, B. duBoulay¹, D. Teather², B. Teather², M. Sharples³ &
G. duBoulay⁴**

¹ School of Cognitive and Computing Sciences, University of Sussex

² Department of Medical Statistics, De Montfort University, Leicester

³ School of Electronic and Electrical Engineering, University of Birmingham

⁴ Institute of Neurology, London

University of Sussex

Cognitive Science Research Paper

CSRP 508

February 1999

ISSN 1350-3162

Abstract: *The nature of expertise has been studied in a wide variety of disciplines. Many of these studies have found that the development of expertise involves qualitative as well as quantitative changes in the cognitive skills and knowledge representations underlying performance. This paper seeks to review the literature which has sought to identify expert/novice differences in the field of diagnostic medical cognition. This literature is concerned with the way experts and novices make diagnostic medical decisions. Differences are identified in terms of hypothesis generation and testing, diagnostic reasoning and the organisation of relevant knowledge. Furthermore, expert's diagnostic reasoning is characterised as largely schema driven, with previous patient encounters influencing medical evaluations of current cases.*

Medical cognition refers to studies of cognitive processes (such as perception, comprehension, decision making and problem solving) in medical practice. Studies have tended to concentrate on the decision making processes and have tended to ignore other areas of medical expertise (such as patient management and reporting), and it should be acknowledged that this gives a rather narrow view of medical expertise. Studies of expertise in diagnostic medical cognition examine differences between practitioners with different levels of experience in terms of their cognitive processes and skills. Many of the studies (e.g. Patel, Arocha & Kaufman, 1994) distinguish between *novices* (individuals who have only everyday knowledge of a domain or the pre-requisite knowledge assumed by the domain, i.e. medical students), *intermediates* (individuals who are above the beginner level but below the sub-expert level, for example, medical residents), *sub-experts* (individuals with generic knowledge but inadequate specialised knowledge of the domain, for example, cardiology experts solving problems in the area of endocrinology) and *experts* (an individual with specialised knowledge of the domain, for example, cardiology experts solving cardiology problems). Some studies have broken these categories down even further, for example, by distinguishing between novices with different levels of experience (*early*, *intermediate* and *advanced novices* referring to 1st, 2nd and 3rd year students respectively - Arocha & Patel, 1995), or between *basic-experts* and *super-experts* (Raufaste, Eyrolle & Marine, 1998). Super-experts refer to the top experts in a particular field.

The aim of this paper is to provide a review of the literature examining expertise effects

on diagnostic medical cognition. Generally speaking, different studies have concentrated on different aspects of expertise effects on cognition, including hypothesis generation and evaluation, memory performance, diagnostic reasoning and the organisation of clinical knowledge. These issues will be dealt with independently in this paper. The paper also examines the more perceptually skilled discipline of radiology. Diagnosis in radiology involves a two step process. First, the radiologists must perform a visual search of the radiograph to examine whether any abnormalities are present in the image. Second, if any abnormalities are identified by the scan, then the radiologist must further examine the areas of abnormality in order to perform a diagnosis. Hence, radiology can be seen as a perceptually skilled discipline. If abnormalities are missed in the first scan of the image, then diagnosis may not be complete and accurate. We may therefore expect radiologists to possess different cognitive skills than experts specialising in domains of medicine with a lesser visual component. This paper does not however concentrate on the visual perceptual skills per se of radiologists - instead the emphasis is on the second stage, the diagnostic process itself, and on the interaction between perception and problem solving.

Most studies addressing the issue of expertise in medical cognition have tended to use the same basic experimental paradigm in which subjects are shown a written description of a clinical case (Patel & Groen, 1986, Patel, Groen & Arocha, 1990) and instructed to read the case notes for a specific period of time. The notes are then removed and the subject asked to recall details of the case and then to provide a diagnosis. In a variation of this paradigm, subjects are presented with the clinical information sequentially, one sentence at a time and asked to explain the incoming information and suggest a diagnosis. These methods are known as the *immediate presentation paradigm* and the *sequential presentation paradigm* respectively. The studies then use semantic, propositional and conceptual analysis techniques to characterise the structure of the subject's knowledge, to examine differences in problem representation and finally to evaluate levels of coherence in the subject's response (see Patel, Arocha & Kaufman (1994) for further detail). Think aloud instructions are also often included. Some researchers have questioned the extent to which these rather artificial experimental tasks allow any valid conclusions to be drawn with regard to the nature of the diagnostic process. Patel, Evans & Kaufman (1989) moved away from the written clinical presentations and instead examined doctor/patient interviews. Medical practitioners were asked to interview patients with the goal of developing a differential diagnosis.

Patients were paid volunteer outpatients who had recently presented with the history and symptoms of an endocrine disorder (the study required the diagnosis of an endocrine disorder). The interviews were conducted in actual settings with realistic time constraints. Physicians were allowed to ask the patient any questions they liked in producing their diagnosis. Although this kind of study can be seen as having more ecological validity than the laboratory based studies discussed earlier, there is still an argument that even this type of study fails to capture the real processes undertaken in producing a diagnosis. Indeed, some researchers (e.g. Klein, Calderwood & MacGregor, 1989, Huber, 1997) believe that in examining real world decisions a more naturalistic approach needs to be taken. Such an approach is based on the belief that real world decisions are influenced by the context in which they are made, studying decisions outside their real world context is not a valid approach. Despite their limitations, laboratory studies of diagnostic medical cognition have produced some interesting results. The aim of this paper is to summarise and discuss these major findings.

Expertise effects

One of the earliest stages in the diagnostic process involves the formulation of working hypotheses. These are derived from clinical observations (in the form of signs, symptoms, test results, and results of physical examinations). The first section of the literature review examines the issue of the effects of expertise on hypothesis generation and evaluation.

1. Hypothesis Generation & Evaluation.

Sisson, Donnelly, Hess and Woolliscroft (1991) examined how experts (with various different areas of sub-speciality within Internal Medicine) and novices (3rd year medical students) differ in terms of the number, specificity and breadth (range) of diagnostic hypotheses generated early in the evaluative process. They argue that these initial hypotheses, although modified by subsequent data, will have a large influence on the diagnostic process. Their results showed that on average, the medical students generated more hypotheses than did the physicians, and this finding was consistent across scenarios. Furthermore, physicians' hypotheses were found to be more general than the students; however, the two groups did not differ in terms of the breadth of their hypotheses (calculated as a percentage of possible categories of diagnosis). Neither group typically named all of the diagnostic categories that logically might have been

included as part of their hypotheses. Sisson et al. assume that the fact that physicians generate less specific hypotheses (in fact, their hypotheses were quite general) may reflect a consciously learned approach or an intuitive evolution in reasoning. The problem with highly specific diagnoses based on limited data is that they may result in some form of premature closure, hence specific hypotheses can be seen as a pitfall that reduces the options available to lead to the proper diagnosis. A final, perhaps surprising finding was that the numbers of hypotheses generated by the individual participants was fairly consistent across the three tasks studied. Other research suggests that the speed with which initial hypotheses are generated is a striking feature of the behaviour of experts. Furthermore, there is evidence that the earlier a good hypothesis set is created, the more predictive it is of the quality of the diagnosis (Joseph & Patel, 1990) .

Lesgold, Rubinson, Feltovich, Glaser, Klopfer & Wang (1988) found that novice (resident) radiologists examining chest x-rays appear to restrict their responses to the most obvious explanation. Two hypotheses were put forward to explain these findings: first, it may be that when there is a dominant hypothesis and a more remote possibility; consideration of the more remote possibility depends upon the availability of mental processing capacity. If any subprocesses of diagnosis are inefficient they will interfere with the more remote response. An alternative explanation is that novices simply do not generate the full range of sensible possibilities in forms that will survive testing and verification. Novices may have learned the triggering rule for the most obvious explanation but not for the subtle special cases. Novices may fail because they have not yet developed the fine-tuned visual acuity needed for feature discrimination that is seen in their more experienced colleagues.

Patel, Arocha & Kaufman (1994), Joseph & Patel (1990) found that experts (endocrinologists solving endocrine problems) produce their hypotheses fairly quickly and accommodate subsequently presented data without introducing any new hypotheses. In contrast, sub-experts (cardiologists solving endocrine problems) continue to generate new hypotheses even after producing most of the diagnostic components needed for the final diagnosis. They are less able to evaluate their hypotheses and hence show an inability to rule out diagnostic hypotheses they had produced earlier. In general, the experts narrowed uncertainty whereas the sub-experts

increased it. They suggest that sub-experts do not have sufficient domain knowledge to discriminate hypotheses. In contrast, experts' initial hypothesis set appears to be particularly well constructed allowing them to complete the diagnostic task from the hypotheses contained in the original set without adding new hypotheses.

Arocha & Patel (1995) examined the effects of inconsistent data on subjects' hypotheses generation and evaluation. Early, intermediate and advanced novices (2nd, 3rd & 4th year medical students) showed differences in terms of their use of co-ordinating operations. These are responses to inconsistent data that are commonly used in scientific and everyday reasoning and include; ignoring data, excluding data, re-interpreting data, re-interpreting hypothesis, modifying a hypothesis to fit the data, and changing a hypothesis altogether. Early and intermediate novices performed more data operations and they more frequently ignored or reinterpreted inconsistent data, whereas advanced novices more often changed their hypotheses to account for the data, changes which decreased the inconsistency with the data. Advanced novices generated a number of early hypotheses and this allowed them to narrow their initial hypothesis set in the face of inconsistent evidence and make fewer data reinterpretations. Intermediates also generated a number of early hypotheses, however they failed to change their hypotheses when confronted with inconsistent data. Instead they maintained several hypotheses of a diverse nature concurrently without evaluating them efficiently. Arocha & Patel concluded that training appears to make novices more sensitive to what their data tells them.

Raufaste, Eyrolle & Marine (1998) were concerned with how the generation of pertinent hypotheses may be related to expertise in radiological diagnosis (experts are known to be particularly adept at producing pertinent hypotheses rapidly). They use a model of spreading activation within semantic networks to explain how experts produce more pertinent diagnostic hypotheses. In particular, pertinence is hypothesised to originate from the interrelations between the elements in the semantic networks - the more schema are linked together, the more pertinence will be demonstrated because pertinent schema will receive multiple primes and then receive more activation. To test this hypothesis the authors examined the semantic networks of novices, intermediates, basic experts and super-experts (top specialists) diagnosing two difficult chest x-rays. As expected, a significant correlation between integration and pertinence was obtained - and

this relationship could not be explained in terms of experience. However it seems that this was not the full story. Whilst intermediates and basic-experts could not be distinguished in terms of the pertinence of their diagnostic hypotheses, super-experts showed a higher level of pertinence than these other two groups, and furthermore the relationship between pertinence and integration disappeared for these subjects. Hence, at this level of performance, an alternative explanation for pertinence must be found. They suggest that in typical cases (the type likely to be seen by intermediates and basic-experts alike), spreading activation would be sufficient to allow the retrieval of pertinent hypotheses. However, for less typical cases, processing would become unconstrained. In such cases, only the super-experts who maintain deliberate reasoning activity in their everyday work (through their research interests and diagnosing unusual cases) would trigger sufficient levels of deliberate reasoning - the kind necessary to constrain possibilities optimally in such cases. Such a hypothesis is consistent with Ericsson, Krampe & Tesch-Romer's (1993) view that 'eminent performance' is directly related to the 'amount of deliberate practice related to that goal'. Interestingly, the authors also found that the richness of the semantic networks underlying problem solving was better able to explain diagnostic success than years of experience.

Another important area of research has examined the effects of expertise on the directionality of reasoning employed. This refers to the distinction between forward reasoning (in which the physician attempts to generate a hypothesis from the findings in a case) and backward reasoning (in which the data collection is influenced by the hypothesis generated). Patel & Groen (1986) found that directionality of reasoning is related to diagnostic accuracy. They found that the diagnostic explanations of subjects making an accurate diagnoses (of acute bacterial endocarditis) consisted of pure forward reasoning. In contrast, subjects with inaccurate diagnoses tended to make use of forward and backward reasoning. Subsequent studies have confirmed this relationship between diagnostic accuracy and forward reasoning. Patel, Groen & Arocha (1990) examined factors which may disrupt this pattern of forward reasoning. They examined the reasoning strategies used by expert cardiologists and endocrinologists solving problems within or outside their area of expertise. They found that an accurate diagnosis (usually from the experts) was associated with pure forward reasoning in the production of an explanation of the principle component of the diagnosis. However, this was often accompanied by one or two components of backward reasoning to explain any loose ends. In contrast, inaccurate diagnoses (usually from the

sub-experts) were associated with the use of both forward and backwards reasoning. According to Patel & Groen (1991), it is important to distinguish between experts' use of backward reasoning (where it is used to tie up loose ends) and that of less experienced physicians. The distinction between forward and backward reasoning is closely related to the distinction between strong-problem solving methods (which are highly constrained by the problem solving environment) and weak methods (which are only minimally constrained). However the two are logically independent. Forward reasoning is highly error prone in the absence of adequate domain knowledge as there are no built-in checks of legitimacy (a great deal of knowledge is needed to use forward reasoning successfully, hence to all intents and purposes it is a strong method). Backward reasoning is best used when domain knowledge is inadequate, as reasoning will be minimally hampered by this lack of knowledge. A weak method is preferable when (as is likely with anyone but an expert) relevant prior knowledge is lacking. Patel & Groen (1991) classify experts' use of backwards reasoning (to tie up loose ends) as a strong method, this is important as the use of backward reasoning is usually associated with the use of weak methods. Lesgold (1988) found evidence that radiological diagnosis is neither primarily top-down (backward reasoning) nor bottom-up (forward reasoning), instead processing is seen as a combination of the two, incorporating features of both models in a recursive, interactive decision making process.

Arocha, Patel & Patel (1993) found differences between novices (medical students) with different levels of expertise. Early novices use a strategy similar to depth first search, considering and evaluating a single hypothesis at a time, furthermore they demonstrate a tendency to maintain hypotheses despite contradictory evidence. Intermediate and advanced novices use a form of breadth first search as they consider and evaluate several hypotheses concurrently. However, differences between intermediates and advanced novices were found, in that intermediates were less skilled at evaluating hypotheses, and hence tended to maintain several hypotheses for long periods of time without resolving or eliminating them. Intermediates also showed a tendency to generate several diagnostic hypotheses to account for different findings. In contrast, advanced novices generated multiple hypotheses to account for the same set of findings. The net result is less irrelevant search compared to the intermediates. Breadth first search has been associated in other areas of reasoning (e.g. scientific reasoning - Klahr & Dunbar, 1988) with problem solving success.

Evans & Gadd (1989) describe four different levels into which clinical knowledge is organised in a medical problem solving context. Observations are units of information that are recognised as potentially relevant in a problem solving context, however they do not constitute clinically useful facts. Findings are observations that have potential clinical significance (e.g. symptoms). Facets are clusters of findings that are suggestive of pre-diagnostic interpretations. They reflect general pathological descriptions such as aortic insufficiency, or categorical descriptions such as endocrine problem. They are also interim hypotheses that divide the information in the problem into manageable sub-problems and suggest possible solutions. Facets vary in their level of abstraction - from high level facets which may partition the problem space and may be a reasonable approximation to a candidate solution - to low level facets which may involve a more local inference that may explain one or two findings and would not advance the problem solving process to the same extent. Diagnosis is the level of classification that encompasses and explains all levels beneath it. The model is hierarchical with facets and diagnoses serving both to establish a context in which observations and findings are interpreted, and to provide a basis for anticipating and searching for confirming or discriminating findings. Arocha & Patel (1995) found that whilst all early and most intermediate novices generated hypotheses at the diagnostic level, most advanced novices generated hypotheses at the facet level. Facet-level hypotheses typically describe general categories of disease with similar underlying processes. Patel, Arocha & Kaufman (1994) found that expert physicians working within their domain of expertise (cardiologists solving cardiology problems) generate hypotheses in the form of high level facets. This serves to partition the problem into manageable units, thus reducing the load on working memory. In contrast, sub-experts (endocrinologists solving cardiology problems) generated hypotheses mostly at the low-level facet, with some high level facets and diagnostic hypotheses.

Patel, Arocha & Kaufman (1994) proposed that a facet can be construed as a retrieval structure that can be used to access rapidly schema from long-term memory (LTM) and to partition a medical problem into manageable units to facilitate the production of a diagnostic hypothesis. This hypothesis has parallels with skilled memory theory, which assumes that at the

time of encoding, experts acquire a set of retrieval cues that are associated in a functional manner with the information to be stored in memory (Ericsson, Krampe & Tesch-Romer, 1993). In a problem solving situation, an expert can use these retrieval structures to provide selective and rapid access to long-term memory.

Another strand of research has examined memory performance as a measure of expertise. Indeed, there is a long tradition expanding over three decades of using memory measures to examine expertise in other disciplines. For example, de Groot (1965) and Chase & Simon (1973) both used recall of mid-game chess positions following brief exposure to the board (approx 5 secs) and found that experts are typically able to recall about 90% of the positions of the pieces accurately, as long as the positions are legal. With random positions the experts do not perform any better than novices. These kinds of findings have been replicated with experts in other fields, including computer programming (Adelson 1984) and bridge (Charness, 1979). Superior memory for the domain is thought to result from differences in the amount of information encoded. Specifically, experts are assumed to encode information in larger, more domain structured chunks than novices.

2. Memory performance as a measure of medical expertise.

Several studies have examined the effects of medical expertise on recall performance with the expectation of demonstrating the usual expert/novice differences. Muzzin, Norman, Jacoby, Freightner, Tugwell & Guyatt (1982, 1983) found that whilst clinicians faced with a written description of a patient's problem processed information more rapidly and in larger chunks than novices, they recalled fewer chunks. Hence there was no overall difference in recall. Groen & Patel (1985) suggested that the failure to find differences between novices and experts may be due to the use of analytic methods which are too simplistic. In their studies they used propositional analysis to examine memory performance. Propositional analysis segments the original text into individual propositions corresponding to the discrete idea units in the text. A similar analysis is carried out on the transcripts recalled by each subject. Hence, propositional analysis is more concerned with the overlap between the underlying meaning of the original and recalled text than with verbatim recall per se. Despite this change in analytic methods, Patel and her colleagues have failed to find any consistent effects of expertise on memory (e.g. Coughlin &

Patel, 1987, Patel, Groen & Fredicson, 1986). To explain these differences from non-medical domains, Norman, Brooks & Allen (1989) point to differences in the nature of the tasks between the highly visual, automatic processing required in chess or bridge, and the much more deliberate, conscious processing of medical information. Patel & Groen (1986) also found no differences between experts (cardiologists solving a cardiology task) and sub-experts (surgeons and psychiatrists solving cardiology problems) in terms of recall. This was true for both relevant and irrelevant propositions. Patel, Arocha & Kaufman (1994) replicated this finding, furthermore, a ceiling effect was found for relevant material. They suggest that recall measures do not discriminate at this level of expertise.

Norman et al. (1989) examined the effects of memory test procedure on expertise effects. In particular, they examined the hypothesis that expertise effects may be more likely to arise when subjects are not aware of the fact that their memory for the cases is to be tested. The incidental memory test condition required subjects to recall case information (laboratory data) after they had provided a diagnosis. They found that the medical students (3rd years) recalled nearly twice as much information in the intentional condition whilst the experts' (nephrologists & respirologists) recall was superior in the incidental condition. Overall, experts recalled more data than the novices, however, this effect was far stronger in the incidental condition. Experts and novices both recalled more critical and abnormal data than non-critical and normal data, however the experts recalled more non-critical data than the novices. They suggest that this may be because having produced a solution, the experts devote more resources to ensuring that this explanation is consistent with the rest of the data present in the case. Experts appear to deal with more pieces of information, as supported by the finding that they mention nearly twice as much data in describing the case. Hence, expertise effects on memory have been demonstrated, however these are stronger when subjects are unaware that their memory is to be tested.

The intermediate effect in recall refers to the finding that intermediates perform better on memory tasks than either experts or novices (Patel & Groen, 1991, Schmidt & Boshuizen, 1993). Patel and Groen also found that whilst intermediates operated on both high and low relevance information showing an inability to discriminate, experts operated more on highly relevant information. Taken together, these findings suggest that outcome measures emphasising recall are

not accurate indices of underlying knowledge and do not measure effective use of knowledge. Recall is nonmonotonically related to expertise, with intermediates recalling more than experts or novices. Schmidt & Boshuizen (1993) found that the intermediate effect in recall disappears if short exposure times are used (30 secs). They suggest that under such restricted conditions, intermediates cannot engage in extraneous search and this affects their recall performance, i.e. they suggest that intermediates process too much 'garbage' - information that experts ignore.

So far, this review has concentrated upon the hypothesis generation, evaluation and memory skills of individuals with different levels of medical expertise. A third area of research has examined diagnostic reasoning skills - the process by which experts and novices draw diagnostic conclusions regarding the nature of a complaint.

3. Diagnostic Reasoning.

Joseph & Patel (1990) found that both experts (endocrinologists solving endocrine cases) and sub-experts (cardiologists solving endocrine cases) selected more relevant and critical cues than irrelevant cues from a case history. However, experts were better able to focus on the critical and relevant information than sub-experts. Experts also generated more links to relate critical or relevant cues, showing better organisation of their domain knowledge. Similar findings have been obtained in radiology. Lesgold et al. (1988) found that experts reported more different findings, had longer reasoning chains, larger and more clusters, and a greater number of their findings connected to at least one other finding. Hence experts differ from novices in terms of the coherence of their knowledge and explanations.

As expected, on most tasks experts show superior diagnostic performance in classifying clinical conditions. Furthermore, there is evidence to suggest that the superior diagnostic skills of experts may be due to their ability to fully utilise patient information presented in clinical protocols (known as enabling conditions). Schmidt, Hobus, Patel & Boshuizen (1987) found that when expert family doctors were presented with slides containing a picture of the patient (allowing them to deduce the subjects age, sex etc), and information about the patients profession, previous diseases, medication, marital status and so forth, then they showed superior diagnostic skills (compared to novice family doctors) when they were subsequently presented

with the patient's complaint (38% vs. 27%). They suggest that this may be due to differential use of enabling conditions information as the experts remembered 40% more of this information. Hobus, Hofstra, Boshuizen & Schmidt (1989) found that when information about enabling conditions was not presented to expert physicians, their diagnostic performance was no better than that of novice physicians. They also found that when experts were asked to describe typical patients with certain diseases, they gave richer descriptions, in particular with reference to the contextual factors facilitating the emergence of the disease.

Similar findings have been obtained with experts in radiology. Norman, Brooks, Coblenz & Babcock (1992) examined the effects of information contained in a brief clinical history on diagnosis and on feature identification from radiographs. They were concerned with the diagnosis of bronchitis, and a bronchitis history (fever, cough and tachypnea) was induced for some trials. Having been presented with the case, paediatric radiologists were asked to rate the likelihood of bronchitis on a scale from -3 for definitely absent to +3 for definitely present. Feature history was found to affect both diagnosis (with a change in the scale of between one half and one unit on the scale), and feature identification (the effect amounted to an increase of about 25%-50% in the number of features identified on film). Furthermore, no novice-expert differences were found in terms of the ratings of likelihood of bronchitis, however, a small significant effect of expertise was found in feature identification with novices reporting more features than experts. They concluded that both experts and novices were equally susceptible to influence from a history. Hence it seems that experience acquired in the normal course of practice is insufficient to avoid such effects.

There is also evidence to suggest that experts and novices differ in their use of biomedical knowledge in making a diagnosis. In particular, findings from studies where subjects are required to think aloud whilst making a diagnosis suggest that experts rarely refer to pathophysiological concepts whilst reasoning about a case, however students use pathophysiological concepts extensively (Boshuizen, Schmidt & Coughlin 1988, Patel, Evans & Groen, 1988). Patel, Groen & Arocha (1994) found that whilst experts (endocrinologists solving endocrine problems) do not use detailed biomedical knowledge in solving routine clinical problems in their own domain of

expertise, they do use it if they are solving either a difficult problem or a problem outside their domain of expertise (cardiologists solving endocrine problems). Boshuizen and Schmidt (1992) examine three possible explanations for these expert/novice differences. Firstly, it may be that after a certain amount of time in the development of expertise, biomedical knowledge becomes rudimentary; with detailed knowledge no longer retrievable. Alternatively, biomedical knowledge may become inert; that is, it is still available to medical experts and can be activated when directly addressed, but is not used in diagnostic reasoning; clinical knowledge is applied instead. The final alternative is that biomedical knowledge may become encapsulated and is integrated in clinical knowledge. According to this hypothesis, experts apply encapsulated knowledge which is associated with detailed, deep level, knowledge which can be retrieved as necessary. Their own evidence supported the encapsulation hypothesis, experts showed increased linkage between concepts used in the think-aloud protocols and the post hoc explanations. They concluded that experts use biomedical knowledge in a tacit way, because in the course of becoming an expert this type of causal knowledge becomes encapsulated into clinical concepts.

According to Gale & Marsden (1983) when students and clinicians are presented with clinical information they interpret it by identifying for themselves personally important pieces of information called 'forceful features'. These act as a key to particular memory structures which in turn give rise to clinical interpretation. They examined experts and novices' diagnostic reasoning by presenting them with cases and asking them to give up to five ideas as to what may be wrong with patients and furthermore, to identify the item of information (forceful feature) which gave rise to these interpretations. They failed to find any differences between experts (registrars and consultants in general medicine) and novices (1st & 3rd year medical students) in terms of the number of interpretations provided, and this finding lead them to conclude that level of clinical experience does not affect the number of interpretations made in response to any given array of clinical information. However, they did find differences in the content of thought. A large amount of variability was found not just between groups but also within them. In every case there was found to be a small focal area of common interpretation and a massive peripheral field of individual differences. No differences in the number of forceful features were identified, but differences in the actual forceful features themselves were found. There was found to be little overlap, both between and within groups. When they examined the forceful features associated

with the correct diagnosis, they found that the more difficult diagnoses allowed identification of far fewer forceful features than the easy diagnoses. They concluded that individuals develop various ways of accessing their memory structures in the case of easy diagnoses, however people do not demonstrate such extensive networks for difficult diagnoses. Highly individualised multiple responses to an array of clinical information are associated with ease of diagnosis. When a diagnosis is being missed, few structures to the right memory structure are found amongst those who think of it and none amongst those who don't. They concluded that experience is not characterised by uniformity of thinking but by individuality of thought. In diagnostic thinking there is no best way, no key piece of information. Experts vary greatly as do students, the difference is that experts have used and changed the knowledge stored so that it becomes progressively more useful personally.

Moskowitz, Kuipers & Kassirer (1988) suggest that the use of heuristics dominates experts clinical problem solving and probabilistic judgments. They distinguish between categorical descriptions of likelihood (e.g. low, moderate, high), ordinal descriptions of likelihood (where quantities are described in terms of ordinal relations, such as greater or less than a reference value) and numerical descriptions of likelihood (where quantities are described in terms of numerical measures). They found that experts references to likelihood consisted primarily of the first two (categorical and ordinal - where numerical values were used, they were used to define categories (70-80%) or as an anchoring value for an ordinal description). Furthermore, the widespread use of two heuristics (representativeness and anchoring) was found. Tied to categorical descriptions of likelihood are archetypic examples, with likelihood categories of new observations often assigned by their closeness of match to these examples. For example, a risk category of very high, might be associated with a small set of examples of very high risk situations. The risk of a new situation is then assessed by determining it's similarity to these examples. This method of risk classification is essentially the representativeness heuristic. Such a short cut often yields valid results because representativeness often correlates with likelihood, however, systematic biases can be introduced as the use of such a heuristic can often result in the neglect of prior probability; be insensitive to a finding's degree of predictability; and to violate the conjunction rule in estimating the likelihood of compound events. Similarly, ordinal descriptions

of likelihood are defined by boundary values, and derived by 'adjusting' away from these anchoring reference points. They suggest that the use of anchoring heuristics can often bias estimates in the direction of the anchoring landmark. According to such a view, the estimation of risk associated with a particular medical procedure will vary depending upon whether it is anchored against the risk of a higher or lower risk procedure. They suggest that the use of heuristics may underlie many clinical logical flaws and explain differences between clinicians in terms of their estimates of uncertainty. Regehr & Norman (1996) also characterise experts' diagnostic decision making as dominated by the use of heuristics. However, they claim that the accuracy of experts suggests that this is not necessarily a bad thing. Furthermore, they argue that instead of teaching learners to avoid heuristics, it may be more useful to help them recognise those few occasions when they are likely to fail. It seems that one of the problems with heuristics is that they may not allow physicians to correctly judge the likelihood of relatively rare conditions. However, Medin & Edelson (1988) found some heuristics are used to allow subjects to take rare cases into account, for example, when confronted with specific cues. Weber, Bockenholt, Hilton & Wallace (1993) found that physicians make a compromise between likelihood and clinical severity because they cannot afford to miss diagnoses with severe consequences.

It is perhaps not surprising given the huge information processing demands of the diagnostic task, that medical decision making has been characterised as dominated by the use of short cuts and heuristics. However, not everyone believes that the use of such shortcuts is an inevitable consequence of the nature of the task. The Centre for Evidence Based Medicine aims to promote evidence based health care by encouraging more controlled evaluative studies and the circulation of results in a format that allows practitioners (given their time constraints) to keep up to date with latest developments. Cochrane (1972) acknowledged that 'it is surely a great criticism of our profession that we have not organized a critical summary, by speciality or sub-speciality, adapted periodically, of all randomised controlled trials.' His observation was that people who want to make more informed decisions about health care do not have ready access to reliable reviews of the available evidence. He called for more systematic reviews of randomised controlled trials (RCT's) in a wide range of medical domains. Reviews of research evidence are necessary to keep practitioners up to date. Such reviews would allow practitioners to make well

informed diagnoses reducing their reliance on heuristics. Such an approach to medical practice is supported by the British Medical Association and the UK government.

A final area of interest in expertise effects in medical cognition is concerned with examining differences between experts and novices in terms of the organisation and structure of their relevant knowledge.

4. Knowledge Organisation.

Arocha & Patel (1995) concluded that the development of expertise requires two phases of learning - rule based learning (i.e. through textbooks and lectures) and experience-based learning (i.e. through exposure to real patients). This is accompanied by knowledge re-organisation, i.e. it involves a qualitative rather than a quantitative change. Joseph & Patel, (1990) found that experts working within their field of expertise (endocrinologists solving endocrine cases) organised information in a coherent form with strong causal relations between various selective data, whereas sub-experts (cardiologists solving endocrine problems) linked information with a greater use of weak conditional and associative relations. This pattern is unlike that of novices, who are unable to select the relevant information from the case (Patel, Green and Frederiksen, 1986). Patel et al. (1994) suggest that one of the differences between intermediates and experts is that the intermediates may not yet have acquired an extensive body of knowledge in a functional manner to perform various tasks. Instead, their knowledge is organised as a flat structure and this results in considerable search, making it difficult for intermediates to set up structures for rapid encoding and selective retrieval of information.

According to the 'small worlds' hypothesis of Kushniruk, Patel & Marley (1998), expert physicians organize diagnostic knowledge on the basis of similarities between disease categories, forming small worlds consisting of small subsets of diseases and their distinguishing features. They speculate that the process of diagnosis involves the physician focusing on relatively small sets of logically related diseases, (i.e. small worlds) and carrying out a limited number of comparisons among these diseases. It is hypothesised that diseases contained within these 'small worlds' would typically share certain overlapping features, and this is the basis for their membership of that particular 'small world'. However, the diseases contained within a 'small

world' differ in terms of the presence or absence of certain other features, allowing the expert to distinguish between the candidate diseases contained within a 'small world'. In order to investigate the 'small worlds' hypothesis, Kushniruk et al. re-analysed the protocols collected in the studies carried out by Joseph & Patel (1990) and Patel, Evans & Kaufman (1989). They were particularly concerned with examining the networks of relationships among the hypotheses and findings generated by experts (endocrinologists solving endocrine cases) and sub-experts (cardiologists solving endocrine cases). The networks produced by the experts were found to contain few elements (i.e. a limited number of hypotheses and findings) which were tightly connected, displaying a high degree of coherence and relatedness. Furthermore, expert physicians quickly focused on those cues and critical findings in a medical case that most clearly distinguish among competing diagnoses in the hypothesis set under consideration. In contrast, the sets of hypotheses generated by sub-experts often contain large numbers of diagnostic hypotheses each belonging to different disease categories. Kushniruk et al. argue that an expert's knowledge is organised in this way because of the limitations of human memory and processing capacity. Furthermore, this organisation of knowledge affects the way experts perform a diagnosis, (i.e. the experts' knowledge organisation and reasoning processes are viewed as being integrally related).

The 'small worlds' hypothesis is closely related to the concept of schema, indeed many theories of expertise assume that experts processing of clinical cases is schema driven. Schemas are defined in cognitive psychology as hypothetical cognitive structures that allow us to call upon our past experience and knowledge in interpreting the present situation. Schemas are thought to play an important role in facilitating the recognition of significant objects within a problem and in enhancing the ability to recognize typical situations. The presence of schemas is used to explain an expert clinician's ability to pay attention to relevant information only and to diagnose diseases rapidly. According to the 'small worlds' hypothesis physicians manage large amounts of information by restructuring their knowledge into small sets of logically related disease schemas. This allows experts to use more efficient discriminatory strategies in order to rule out competing hypotheses, allowing them to focus on the few critical findings that clearly differentiate between competing hypotheses.

Radiological screening is also thought to be schema driven. Hillard, Myles-Worsley, Johnston & Baxter (1985) proposed that radiologists interpret chest radiographs by means of an internal visual framework that develops through experience and training. As radiologists gain experience in the reading of chest radiographs they develop schema. Normal posterior/anterior (PA) chest radiographs have certain features in common that make up the schema; what is included in the schema is what is normal. Abnormal findings represent schema discrepant information. They tested the hypothesis that experts process normal images automatically whilst abnormal images require further processing, by examining recognition memory for normal and abnormal images. They hypothesised that because the abnormal images represented schema discrepant information, they would be processed further and remembered better than normal images. No such effects were predicted for the non-experts as no processing differences were expected. Consistent with this hypothesis, experts had superior memory for abnormal radiographs but inferior memory for normal radiographs. Myles-Worsley, Johnston & Simons (1988) carried out a similar study and found similar results. There is however a point of clarification: radiological experience does not render observers sensitive to arbitrary deviations from normality, instead experience sensitizes observers only to clinically significant deviations (i.e. those that indicate disease). Radiological expertise appears to reduce sensitivity to deviations that are not clinically significant. According to Lesgold et al. (1988) radiological schema serve at least two functions, firstly the assignment of x-ray features to normal-anatomy schema largely determines which features are left over and hence show signs of possible abnormality. Furthermore, the normal-anatomy schema may contain attachment procedures, or localisation rules, for determining where the abnormality lies. Radiologists are taught to look for a variety of localisation cues which allow them to map abnormalities onto anatomic structures. Lesgold (1984) found that the non-experts' use of localization cues was more deliberate, more fragmentary, and less appropriate.

Schmidt, Norman & Boshuizen's (1990) produced a theory of expertise which, consistent with schema theories, assumes that expertise is not so much a matter of superior reasoning skills or depth of knowledge of pathophysiological states, but rather is based on cognitive structures that describe the features of prototypical or even actual patients. Their model of medical expertise makes several assumptions. First, the model assumes that in acquiring expertise,

students progress through several transitory stages, characterised by distinctively different knowledge structures underlying their performance. Second, they assume that these representations do not decay or become inert in the course of developing expertise but remain available for future use, when the situation requires activation. Finally, the model assumes that experienced physicians, whilst diagnosing routine cases, use knowledge structures or 'illness scripts'. These structures develop through continuous exposure to patients. They contain little knowledge of pathophysiological causes of symptoms and complaints - but a wealth of clinically relevant information about the disease, its consequences and the context in which it develops. There are various levels of illness scripts from categories of diseases to individual patients seen before.

According to this 'stage' model, early in the development of medical knowledge, students develop rich, elaborated causal networks explaining the causes and consequences of disease in terms of general underlying pathophysiological processes. Early learning is based largely on books, and this results in a prototypical perspective on disease, with limited understanding of variability of the disease in reality. At the second stage, brought about through extensive and repeated exposure to real patients, the declarative knowledge outlined in stage 1 becomes compiled into high-level, simplified causal models explaining signs and symptoms, which are subsumed under diagnostic labels. Short cuts become available through experience. Students at this level of development will not have to activate all relevant knowledge in order to understand the patient, only knowledge pertinent to the case. In short, knowledge is re-organised to ensure accessibility and efficient use, into simplified causal models explaining signs and symptoms.

The third stage is reached once the student has had sufficient experience meeting 'real' patients. As a result of meeting patients, the student gets a feeling of how disease manifestations may vary, paying attention to contextual factors under which disease emerges. Instead of causal factors, the different features that characterise the clinical appearance of a disease becomes the anchor point around which their thinking evolves. The student develops a series of representations called 'Illness Scripts'. These structures contain varied information including: enabling conditions - factors making occurrence of a disease more likely (e.g. hereditary factors), including predisposing factors (e.g. drugs) and boundary conditions (e.g. age, sex), faults - a

description of the malfunction, and consequences - the signs and symptoms arising from the fault. Furthermore, illness scripts provide the rules enabling one to construct mental models of a family of diseases (based on feature overlap - in this way the model has parallels with the small worlds hypothesis), a specific disease or even a patient having the disease. Illness scripts also seem to specify a specific order of information - they obey certain conventions regarding an optimal structure in medicine. Coughlin & Patel (1987) found that experts' (family practitioners) memory for case information was more sensitive than novices (2nd year medical students) to the ordering of the data presented. They suggest that experience results in experts expecting information to be presented in a certain order, if the information presented is not consistent with such expectations, then this affects their ability to process the information. Illness scripts, because they develop through experience are highly idiosyncratic in nature. For any disease, an individual physicians scripts may or may not resemble the scripts of other physicians or the textbook. The absence of pathophysiological information in the scripts (apart from a simple model explaining the causes) implies this information is not normally used by the physician. Indeed as discussed earlier, there is evidence that experts do not use this form of data in their diagnostic reasoning.

Finally, having reached the final stage, experts begin to store patient encounters as 'Instance Scripts'. According to this hypothesis, memory for previous patient encounters are retained in memory as individual entities and are not merged into some prototypical form. Furthermore, an assumption is made that, to a large extent, expert clinical reasoning is based on the similarity between the presenting situation and some previous patient available from memory. Hassebrock & Pretula (1990) found that physicians retained vivid autobiographical memory for cases seen as long as 20 years earlier. Van Rossum & Bender (1990) showed the impact of a single vivid case on diagnosis of similar examples two years later. Pattern recognition plays an important part in diagnosis.

Central to the model is the assumption that as clinicians gain experience and develop new information processing structures, previously acquired knowledge remains available and expert physicians may move from one stage to another as the complexity of the problem demands. The model does not suggest that experts work at a deeper level of processing, but rather, that

expertise is associated with the availability of knowledge representations in various forms, derived from both experience and formal education.

Schmidt & Boshuizen (1993) describe an elaborated model of expertise development in which the learner is seen as progressing through a series of consecutive phases, each of which is characterized by functionally different knowledge structures underlying performance. The first phase is characterised by the accumulation of causal networks explaining the causes and consequences of disease in terms of general underlying biological or pathophysiological processes. Through experience with real cases, this knowledge transforms into narrative structures called 'illness scripts'. The cognitive mechanisms responsible for this transition are encapsulation of elaborated knowledge into high level but simplified causal models or even diagnostic categories, and tuning through the inclusion of contextual information. Illness scripts thus contain the physician's encapsulated pathophysiological knowledge of the disease and its consequences, in addition to clinical knowledge of the constraints under which a disease occurs. When solving a problem, a physician searches for an appropriate script. Once one (or a few) have been selected, they will tend to match its elements to the information provided by the patient. Once the script has been instantiated, they remain available in memory as episodic traces of previously diagnosed patients and are used in the diagnosis of future patients. The third stage is characterised by the use of these episodic memories of actual patients in the diagnosis of new cases. They postulated that each type of knowledge forms a layer in memory, and although not usually applied any more, remains available for use when ontologically more recently acquired structures fail to produce an adequate representation of a clinical problem.

Rogers (1995) provided insights into the cognitive skills utilised in diagnostic radiology. She was particularly concerned with describing the interaction between perception and problem solving. According to the model, radiologists' prior knowledge of both the anatomical region under consideration and the imaging modality used, allows the expert to retrieve a particular collection of declarative and procedural knowledge from memory (for example, the anatomical region implies a certain set of anatomical objects, while the modality calls into play knowledge of the kinds of perceptual cues likely to be present). This process typically results in the formulation of a check list allowing the radiologist to conduct a relatively ordered examination. Hence, context

creates expectations as to what the practitioner is likely to see, as a result plans to explore these expectations emerge that then guide the attention process in deliberate search. In radiology, Rogers found that expectations are largely perceptual in nature, especially in the absence of other information about the patient (e.g. test results or physical examination). She also noted that processing is not just top-down, instead there are often unexpected phenomena in the image, which seem to capture attention immediately, and cause currently active plans to be interrupted or abandoned in favour of new exploratory activity. Rogers noted at least two different types of attentional activity. The first is characterised by relatively fast noticing and labelling of an abnormality as soon as the x-ray image appeared. This was termed 'immediate visual capture' and was quite often accompanied by a brief description of the abnormality (e.g. shape and size). The second, deliberate landmark search, requires attention to be focused purposefully and serially. Deliberate landmark search maybe particularly useful when there is ambiguity in the x-ray image which may be removed by a more careful examination of the landmarks.

Rogers was also concerned with the interaction between perception and problem solving. Three different levels of oversights in the transition were identified. At the perceptual level, a detection oversight occurs when the subject fails to see the abnormality, whereas a labelling error occurs when the subject sees the abnormality but fails to correctly label it. Finally, at the problem solving level, integration errors are found in cases where the subject saw and labelled the abnormality correctly, but failed to use this information in the generation of a diagnostic hypotheses.

Discussion

This paper has attempted to provide a review of the literature examining expertise effects in medical cognition. The major areas of investigation for this research have included hypothesis generation and evaluation, recall, diagnostic reasoning and knowledge organisation. The evidence discussed suggests that experts and novices differ in terms of a number of key skills which are related to their diagnostic success. With regard to hypothesis generation and evaluation, the literature suggests that experts produce fewer, but more general hypothesis (at the level of the facet rather than at the level of diagnoses) at an earlier stage of problem formulation than novices.

Furthermore, experts work from findings to a hypothesis (forward reasoning) using a breadth first approach (considering and evaluating several hypothesis at once). In contrast, novices reasoning is characterised as backwards (from hypothesis to data), and furthermore, depth first (considering and evaluating a single hypothesis at a time). Experts also demonstrate superior hypothesis evaluation skills, in particular, they are better able to disregard discredited hypotheses and are more likely to change their hypothesis to fit the data than to change the data to fit their hypothesis or to ignore inconsistent findings altogether.

The evidence suggests that memory measures are a less reliable measure of expertise. Specifically, unlike diagnostic reasoning, recall is nonmonotonically related to expertise. As a general rule, intermediates perform better on memory tasks than either experts or novices, and it is suggested that this is because they process more of the case material that is not relevant to the diagnosis. This processing results in superior memory for a larger number of case details, however because this information is not relevant to the diagnosis, it is not related to diagnostic success. Experts and sub-experts cannot be distinguished in terms of memory measures.

In terms of diagnostic reasoning, experts have been found to make extensive use of information contained within the case notes (known as enabling conditions), furthermore this finding has been found to extend to both general family practitioners (Schmidt, Hobus, Patel & Boshuizen, 1987) and to the more specialised discipline of radiology (Norman, Brooks, Coblenz & Babcook, 1992). However, experts make less use of biomedical terms and explanations in their diagnosis, instead their explanations are based on clinically relevant information. Clinically relevant information is specified by experts' own representation of the disease known as an 'illness script'. 'Forceful features' refer to those personally important pieces of information which give rise to a clinical interpretation. Little consistency has been found in terms of the forceful features associated with a diagnosis by experts or novices. Hence, it seems that experience is not characterised by uniformity of thinking but by individuality of thought, experts vary greatly as do students, the difference is that experts, through experience, have used and changed the knowledge stored so that it becomes progressively more useful personally.

Experts appear to organise their clinical knowledge of different disease in terms of small groups of confusable diseases (known as 'small worlds'). These groups have in common certain

important features which is the basis for their membership of any particular small world. These structures also contain information which allows the clinician to distinguish between the competing diagnostic hypotheses contained within that small world, thus aiding correct diagnosis. Experts have also been found to retain vivid memories of previous patient encounters, in the form of 'instance scripts'. Previous cases diagnosed by an expert have been found to affect diagnosis of a similar case upto 20 years later, with clinical reasoning based on the similarity between the two cases.

Despite the fact that studies of diagnostic medical cognition have shown clear differences in the cognitive skills and knowledge organisation possessed by experts and novices, Johnson (1988) points out that research which has focused upon performance presents a rather pessimistic appraisal of experts, with experts often failing to perform any better than novices. For example, Goldberg (1959) found no differences between psychiatrists and their secretaries in terms of their ability to diagnose brain damage using a common test. Furthermore, Goldberg (1968) found that a regression model performed with greater accuracy than experts in the diagnosis of psychosis using the MMPI (Minnesota Multiphasic Personality Inventory). Johnson (1988) examined physicians' ability to evaluate applicants for internships. He found that a simple regression model was more effective in selecting successful candidates than the experts. To explain these findings, he point out that experts often include in their analysis individual case data that would not be included in a regression model, in other words, experts seem to pay attention to relatively rare variables, applicable only to the case under consideration (these are termed 'broken leg cues'). However, experts are less able to combine more mundane information available for every case (the information used by regression models). This distinction is analogous to the distinction between base-rate and case-specific information, with the base-rate information often underweighted compared to the case-specific information. Regression models are successful because of their ability to properly weight base-rate data. Experts are less able to evaluate such large quantities of data.

Dawes (1988) suggests that regression equations do so well because most of the clinical tasks undertaken are ones in which, 'as far as current knowledge permits us to know', the relationship between the cues we know to have some relevance and the thing we are predicting is

'conditionally monotonic - a higher value on each of these cues is associated with a higher value on the variable being predicted, irrespective of the value of the other cues'. Dawes notes that this is exactly the statistical situation in which regression models are very good predictors. However, this does not mean that the clinician is redundant, rather, they are left with the tasks that we know relatively little about, relative to their immense complexity. The consolation to the clinician is that it is only they who can establish the relevant set of non-redundant cues for the computer to work on.

Johnson's (1988) characterisation of expert's decision processes as dominated by biases and heuristics (in the form of so called 'broken leg cues') is consistent with that outlined by other researchers examining the more cognitive aspects of experts performance. For example, we discussed earlier how experts, when judging likelihood, make widespread use of two heuristics, the representativeness heuristic and the anchoring heuristic (Moskowitz et al, 1988). Furthermore, Johnson's (1988) characterisation of experts as being particularly influenced by case-specific data is consistent with the findings of Schmidt et al. (1988), who characterised experts as being particularly influenced by the patient information contained in the clinical protocols (the enabling conditions). Hence, there is further evidence to support Johnson's claims.

Studies like Johnson's which examine error rates of experts and statistical models have however, been criticised as they fail to take account of the type and consequences of different types of error. Hence it may be that whilst statistical models make fewer errors than experts, the errors the models do make may prove to be more serious (i.e. there may be graver consequences for the patient). This would then prove to be a mitigating factor against the use of statistical models in diagnosis.

A related issue affecting practitioners' ability to produce an accurate diagnosis concerns their use of base rate probabilities. In particular, we can distinguish between two probabilities relevant to diagnosis, firstly, the probability of disease from base rates, and secondly, the probability of disease from the clinical presentation of the case. Accurate diagnosis will be dependent upon a physicians ability to incorporate case data accurately and modify diagnostic probabilities accordingly. There is also a wider issue concerning whether practitioners should use

base rates at all - or simply consider the probability of disease from the presentation of the case. Perhaps the strongest argument that can be made in favour of ignoring base rates is that they penalise rare complaints in favour of more common illnesses.

To conclude, studies of diagnostic medical cognition suggest that experts differ from novices and sub-experts in terms of a number of skills which affect their diagnostic accuracy. In particular, experts are proposed to have better or more complete representations of the task domain, better organisation of task relevant knowledge and superior diagnostic reasoning skills. However, it seems that such a characterisation of experts may not tell the whole story. It seems that experts' reasoning is characterised by the widespread use of various heuristics and influenced by biased data evaluations. The finding that experts tend to underweigh base rate information compared to case-specific information should be approached with caution however, as the task employed by Johnson (1988) was not strictly speaking a medical one. However, the conclusion is given some credibility by the studies discussed earlier in this report which suggest that expert physicians' diagnostic reasoning is strongly influenced by the enabling conditions presented in the clinical data. As pointed out by Johnson (1988), the implications of these findings are of particular relevance to practitioners involved in the design of decision support tools. It is clear that these systems should examine the information that experts currently underweigh and provide a measure of its impact. The expert can then adjust this initial estimate to account for information not considered by the model, such as the enabling conditions, hence taking advantage of the computational abilities of the computer and the skills and experience of expert physicians.

References

Adelson, B. (1984). When novices surpass experts: The difficulty of a task may increase with expertise. Journal of Experimental Psychology, 10, 483-495.

Arocha, J.F., & Patel, V.L. (1995). Novice diagnostic reasoning in medicine: Accounting for Evidence. The Journal of the Learning Sciences, 4(4), 355-384.

Arocha, J.F., Patel, V.L., & Patel, Y.C. (1993) Hypothesis generation and the coordination of

theory and evidence in novice diagnostic reasoning. Medical Decision Making, 13, 198-213.

Boshuizen, H.P.A., Schmidt, H.G., & Coughlin, L.D. (1988). On the application of medical basic science knowledge in clinical reasoning: Implications for structural knowledge differences between experts and novices. Proceedings of the 10th Conference of the Cognitive Science Society, 517-523. Hillsdale, New Jersey: Erlbaum.

Boshuizen, H.P.A. & Schmidt, H.G. (1992). The role of biomedical knowledge in clinical reasoning by experts, intermediates and novices. Cognitive Science, 16, 153-184

Brooks, L.R., Norman, G.R & Allen, S.W. (1991). Role of specific similarity in a medical diagnostic task. Journal of Experimental Psychology: General, 120, 278-287.

Charness, N. (1979). Components of skill in bridge. Canadian Journal of Psychology, 33, 1-16.

Chase, W.G., & Simon, H.A. (1973). Perception in chess. Cognitive Psychology, 1, 55-81.

Cochran, A.L. (1972). Effectiveness and Efficiency. Random Reflections on Health Services. London: Nuffield Provincial Hospitals Trust. (Reprinted in 1989 in association with the British Medical Journal).

Coughlin, L., & Patel, V.L. (1987). Processing of critical information by physicians and medical students. Journal of Medical Education, 62, 818-828.

Dawes, R.M. (1988). You can't systematize human judgment: Dyslexia. In J. Dowie & A. Elstein, Professional Judgment: A Reader in Clinical Decision Making. Cambridge University Press.

de Groot, A.D. (1965). Thought and choice in chess. The Hague: Mouton.

Ericsson, K.A., Krampe, R. & Tesch-Romer, C. (1993) The role of deliberate practice in the acquisition of expert performance, Psychological Review, 100 (3), 363-406.

Evans, D.A., & Gadd, C.S. (1989). Managing coherence and context in medical problem-solving discourse. In D. Evans & V. Patel (Eds.), Cognitive Science in Medicine: Biomedical Modelling, 211-255. MIT Press, Cambridge, Massachusetts.

Gale, J. & Marsden, P. (1983) Medical Diagnosis: From Student to Clinician. Oxford University Press, Oxford.

Goldberg, L.R. (1959). The effectiveness of clinicians' judgments: The diagnosis of organic brain damage from the Bender-Gestalt test. Journal of Consulting Psychology, 23, 25-33.

Goldberg, L.R. (1968). Simple or simple processes? Some research on clinical judgments. American Psychologist, 23, 483-496.

Groen, G. J. & Patel, V. M. (1985). Medical problem-solving: Some questionable assumptions. Medical Education, 19, 95-100.

Hassebrock, F. & Pretula, M. (1990). Autobiographical Memory in medical problem solving. Paper presented at the American Educational Research Association Meeting, Boston, Massachusetts.

Hillard, A., Myles-Worsley M., Johnson, W. & Baxter, B. (1985). The development of Radiologic schemata through training and experience: A preliminary communication. Investigative Radiology, 18 (4), 422-425.

Hobus, P.P.M., Hofstra, M.L., Boshuizen, H.P.A., & Schmidt H.G. (1989). Mental representation of prototypical patients: Expert-Novice differences. Paper presented at the First European Congress on Psychology, Amsterdam.

Huber, O., Wider, R., & Huber, O.W. (1997). Active information search and complete information presentation in naturalistic risky decision tasks. Acta Psychologica, 95(1), 15-29.

Johnson, E.J. (1988). Expertise and decision under uncertainty: Performance and process. In Chi, M., Glaser, R., & Farr, M.J. (Eds.). The Nature of Expertise, 209-228, Hillsdale, NJ: Erlbaum.

Joseph, G.M., & Patel, V.L. (1990). Domain knowledge and hypothesis generation in diagnostic reasoning. Medical Decision Making, 10, 31-46.

Klahr, D., & Dunbar, K. (1988). Dual space search during scientific reasoning. Cognitive Science, 12, 1-55.

Klein, G., Calderwood, R., & MacGregor, D. (1989). Critical decision method for eliciting knowledge, IEEE Trans. Syst., Man Cybern. 19(3), 462-472.

Kushniruk, A.W., Patel, V.L. & Marley, A.J. (1998). Small worlds and medical expertise: implications for medical cognition and knowledge engineering. International Journal of Medical Informatics, 49, 255-271.

Lesgold, A. (1984). Acquiring expertise. In J.R. Anderson & S.M. Kosslyn (Eds.), Tutorials in learning and memory: Essays in honour of Gordon Bower (pp.31-60). San Francisco: W.H. Freeman.

Lesgold, A. (1988). Problem solving. In R.J. Sternberg & E.E. Smith (Eds.), The psychology of human thought, 188-214. Cambridge, MA: Cambridge University Press.

Lesgold, A., Robinson, H., Feltovich, P., Glaser, R., Klopfer, D., & Wang, Y. (1988). Expertise in a complex skill: Diagnosing X-ray pictures. In M. Chi, R. Glaser & M.J. Farr (Eds.), The nature of expertise, 311-342. Hillsdale, NJ: Erlbaum.

Medin, D., & Edelson, S. (1988). Problem structure and the use of base-rate information from

experience. Journal of Experimental Psychology: General, 117(1), 68-85.

Moskowitz, A.J., Kuipers, B.J., & Kassirer, J.P. (1988). Dealing with Uncertainty, Risks, and Tradeoffs in Clinical Decisions: A Cognitive Science Approach. Annals of Internal Medicine, 108, 435-449.

Muzzin, L.J., Norman, G.R., Jacoby, L.L., Freightner, J.W., Tugwell, P., & Guyatt, G.H. (1983). Expertise in recall of clinical protocols in two speciality areas. Proceedings of the 22nd Conference on Research in Medical Education, 122-127. Washington: American Association of Medical Colleges.

Muzzin, L.J., Norman, G.R., Jacoby, L.L., Freightner, J.W., Tugwell, P., & Guyatt, G.H. (1982). Manifestations of expertise in recall of clinical protocols. In Proceedings of the 21st Annual Conference on Research in Medical Education, 163-168. Washington, DC: American Association of Medical Colleges.

Myles-Worsley, M., Johnston, W. A. & Simons, M.A. (1988). The influence of expertise on x-ray image processing. Journal of Experimental Psychology: Learning, Memory & Cognition, 14 (3), 553-557.

Norman, G.R., Brooks, L.R. & Allen, S.W. (1989). Recall by expert medical practitioners as a record of processing attention. Journal of Experimental Psychology: Learning, Memory & Cognition, 15, 1166-1174.

Norman, G.R., Brooks, L.R., Coblenz, C.L. & Babcock, C.J. (1992). The correlation of feature identification and category judgments in diagnostic radiology. Memory & Cognition, 20 (4), 344-355.

Patel, V.L., Arocha, J.F., & Kaufman, D.R. (1994). Diagnostic reasoning and medical expertise. In D. Medin (Ed.), The psychology of learning and motivation (Vol.31), (pp.187-252). San Diego,

CA: Academic Press.

Patel, V.L. & Groen, G.J. (1991). The general and specific nature of medical expertise: A critical look. In K.A. Ericsson & J. Smith (Eds.), Toward a general theory of expertise: Prospects and limits (pp. 93-125). New York: Cambridge University Press.

Patel, V.L., Groen, G.J., & Fredicson, C.H. (1986). Differences between students and physicians in memory for clinical cases. Medical Education, 20, 3-9.

Patel, V.L., Evans, D.A., & Groen, G.J. (1989). Biomedical knowledge in clinical reasoning. In D. Evans & V. Patel (Eds.), Cognitive Science in Medicine: Biomedical Modelling, (pp. 53-121), MIT Press, Cambridge, Massachusetts.

Patel, V.L., Evans, D., & Kaufman, D. (1989). A cognitive framework for doctor-patient interaction, in: D. Evans, V. Patel (Eds.), Cognitive Science in Medicine: Biomedical Modelling, MIT Press, Cambridge, MA, pp. 253-307.

Patel, V.L., Groen, G.J., & Arocha, J.F. (1990). Medical expertise as a function of task difficulty. Memory & Cognition, 18(4), 394-406.

Patel, V.L. & Groen, G.J. (1986). Knowledge-based solution strategies in medical reasoning. Cognitive Science, 10, 91-116.

Raufaste, E., Eyrolle, H & Marine, C. (1998). Pertinence Generation in Radiological Diagnosis: Spreading Activation and the Nature of Expertise, Cognitive Science, 22(4), 517-546.

Regehr, G. & Norman, G.R. (1996). Issues in Cognitive Psychology: Implications for Professional Education. Academic Medicine, 71(9), 998-1001.

Rogers, E. (1995). A Cognitive Theory of Visual Interaction. In B. Chandrasekaran, J. Glasgow & N. Hari Narayanan (Eds.), Diagrammatic Reasoning: Cognitive and Computational Perspectives, AAAI Press/ The MIT Press, Menlo Park, California.

Schmidt, H.G., Hobus, P.P., Patel, V.L. & Boshuizen, H.P.A. (1987). Contextual factors in the activation of first hypotheses: Expert-Novices differences. Paper presented at the Annual Meeting of the American Educational Research Association, Washington, D.C.

Schmidt, H., Norman, G., & Boshuizen, H.P.A. (1990). A cognitive perspective on medical expertise: Theory and implications. Academic Medicine, 65(10), 611-621.

Schmidt, H.G., & Boshuizen, H.P.A. (1993). On the origin of the intermediate effect in clinical case recall. Memory & Cognition, 21, 338-351.

Schmidt, H.G., & Boshuizen, H.P.A. (1993). On Acquiring Expertise in Medicine, Educational Psychology Review, 5, 205-221.

Sisson, J.C., Donnelly, M.B., Hess, G.E., & Woolliscroft, J.O. (1991). The characterisation of early diagnostic hypotheses generated by physicians (experts) and students (novices) at one medical school. Academic Medicine, 66, 607-612.

Van Rossum, H.J.M., & Bender, W.W. (1990). What can be learnt from a boy with Acute Appendicitis? Persistent effects of a case presentation on the diagnostic judgement of family doctors. Paper presented at the Fourth Ottawa Conference, Ottawa, Ontario, Canada.

Weber, E.U., Bockenholt, U., Hilton, D.J., & Wallace, B. (1993). Determinants of diagnostic hypothesis generation: Effects of information, base rates and experience. Journal of Experimental Psychology: Learning, Memory & Cognition, 19(5), 1151-1164.