

# The Development of Human Expertise in a Complex Environment

Michael Harré · Terry Bossomaier · Allan Snyder

Received: 13 October 2010 / Accepted: 21 March 2011 / Published online: 9 April 2011  
© Springer Science+Business Media B.V. 2011

**Abstract** We introduce an innovative technique that quantifies human expertise development in such a way that humans and artificial systems can be directly compared. Using this technique we are able to highlight certain fundamental difficulties associated with the learning of a complex task that humans are still exceptionally better at than their computer counterparts. We demonstrate that expertise goes through significant developmental transitions that have previously been predicted but never explicated. The first signals the onset of a steady increase in global awareness that begins surprisingly late in expertise acquisition. The second transition, reached by only a very few experts in the world, shows a major reorganisation of global contextual knowledge resulting in a relatively minor gain in skill. We are able to show that these empirical findings have consequences for our understanding of the way in which expertise acquisition may be modelled by learning in artificial intelligence systems. This point is emphasised with a novel theoretical result showing explicitly how our findings imply a non-trivial hurdle for learning for suitably complex tasks.

**Keywords** Expertise · Information theory · Games · Decision theory

## Introduction

The human ability to make rapid and accurate decisions in situations where most information needs to be disregarded and only a relatively small amount of

---

This work was supported by ARC grant DP0881829 and by US Air Force grant AOARD 09094.

---

M. Harré (✉) · A. Snyder  
Centre for the Mind, The University of Sydney, Sydney, NSW, Australia  
e-mail: michael.harre@sydney.edu.au

T. Bossomaier  
Centre for Research in Complex Systems, Charles Sturt University, Bathurst, NSW, Australia

information needs to be attended to is unparalleled in artificial systems to date. Obtaining a more thorough understanding of how such human expertise develops in complex environments is one of the fundamental outstanding questions in the cognitive sciences (Ericsson and Charness 1994) as well as one of the key missing ingredients for artificial intelligence (AI) (Dreyfus 2007). There is a long history of vigorous debate between those that claim AI can suitably replicate human thought and those that believe human thought is not computational in a fundamental sense, see for example the views of Simon (1973) and Dreyfus (1992). This paper sheds light on what information is relevant in *populations* of players of the challenging game of Go and the implications this has for our understanding of human expertise and how this is related to the prospects for AI research.

A major view of human expertise from Simon, Gobet, Ericsson and other psychologists (Ross 2006; De Groot and Gobet 1996; Simon 1991; Gobet 2005) is that it operates by the building up of a very large library of patterns and actions conditioned upon them. Gobet (1998) has estimated that anywhere between 10,000 and 100,000 of these patterns need to be stored in long term memory in order for a player to become an expert in their discipline. Gobet and Simon (1996) and Ericsson et al. (1993) have noted that this takes in the order of 10 years of deliberate practice to achieve.

Expertise in games has long been an important area of study in both psychology and artificial intelligence. The prodigious talents of expert Chess players for example have often intrigued laymen and psychologists alike. An important focus of research in this area has been to discover the limitations that exist for our cognitive abilities [see for example De Groot and Gobet (1996), Reitman (1976)]. On the other hand, progress in replicating the impressive abilities humans have in complex tasks by using computers has been surprisingly slow (Brooks 1990; Gamez 2008; Wang 2007; Ekbja 2008). For example Deep Blue, the chess playing IBM computer, was much stronger in terms of brute force computational ability (calculations per second, memory capacity and breadth of search for example) but was still only just able to beat Kasparov (Campbell et al. 2002).

The empirical studies of human thinking and decision-making have revealed a number of general shortcomings (Kuo et al. 2008) which have been modelled (Engle et al. 1999; Newman and Polk 2008) and ways of improving them have been found (Jaeggi et al. 2008; Buschkuhl et al. 2008). Simply defined problems, requiring short solutions (which may be very difficult as in advanced cognitive ability assessment or working memory tests) are often performed by large numbers of individuals. But complex decision making is rarely studied on such a scale. Typical neuro-economic experiments for example would rarely have cohorts much above 20 [i.e. see Deppe et al. (2005)]. The situation in which brain imaging is involved reduces the number of subjects even further. Thus the outcomes have a strong individual character and do not necessarily generalise to *collective human expertise* in any domain. Our study precisely determines the properties of expert decision making across whole *populations*.

We circumvent the small cohort difficulty by choosing to study Go. This game is currently by far the most difficult game for computers amongst all established games. Thus Go expertise is *a crucially important indicator of the unique factors of*

*human performance*. It is not that Go is yet to be solved by computers, but rather that its complexity is sufficiently large to model what humans confront and we were able to build move by move data from 160,000 players—from skilled amateurs to the elite professionals of whom there are only a few hundred worldwide. These players have the ultimate rank of *9Dan Professional*. We refer to them as the *grand-masters* although, unlike Chess, this term is not used amongst Go players. Consequently, we can compare every stage of development to that of a grand-master confronting the same situations. Our study reveals two unexpected transitions in the progress from amateur to world champion rank.

## The Game of Go

The game of Go has played a significant cultural role in parts of Asia for more than 2,500 years.<sup>1</sup> Historically Go originated in China but it is now widely played in Japan and Korea and has recently become popular in many western countries. There are international competitions in which professional players can earn considerable money as well as tournaments that have developed for the sole purpose of testing the ability of AI systems to play one another.

The rules of Go are remarkably simple. There are two players, one plays with white pieces, called *stones*, and the other with black. The board on which the game is played is most commonly a grid  $19 \times 19$  in size. The two players take it in turn in placing their stones on any unoccupied intersection of two gridlines on the board. Each stone has *liberties*, immediately adjacent positions that are unoccupied by another stone (the maximum number of liberties for a single stone is therefore 4). Stones of the same colour can also form chains: if a stone occupies the liberty of another stone, then they are said to form a chain or group and they share their joint liberties, i.e. two stones have a maximum of six shared liberties. Chains can be arbitrarily long and the whole chain is considered as one entity with possibly a large number of liberties. Once placed on the board a stone cannot be moved unless it is captured. A capture occurs when a stone or a group of stones have all of their liberties occupied by stones of another colour. In this case all of the stones in the group are captured and removed from the board.

The aim of the game is to completely surround as much territory as possible, the player with the most territory surrounded wins the game. In terms of combinatorial game theory, Go is a deterministic, two player, zero sum, perfect information game. It is considerably more complex than Chess in terms of both state space complexity and game tree complexity (Tromp and Farneback 2006) and has so far eluded brute force search techniques for finding good moves, as Deep Blue successfully managed for Chess, although some recent progress has been made (Wang and Gelly 2007).

Go has not played as prominent a role in psychological studies as Chess has but there have been some interesting results to come from the limited literature available. A comparison of fMRI studies for Chess (Atherton et al. 2003; Go Chen et al. 2003) has shown that there are strong similarities as well as qualitative

---

<sup>1</sup> There is a great deal of information regarding Go available at: <http://senseis.xmp.net>.

differences in the regions of the brain that are activated by Chess players versus Go players. Of particular note is that Go appears to activate regions associated with semantic decision tasks whereas Chess does not (or at least not to the same extent). In an earlier study, Reitman (1976) used inter-response times in a Go board reconstruction task to show that Go (sub-)positions are not linearly segregated into chunks in the mind of expert players (as they are in Chess (Chase and Simon 1973), but instead are overlapping clusters of pieces, usually two or three stones in each, that interact in ways quite different to that found for Chess pieces. Despite this apparent difference, Stern et al. (2006) have shown that predicting the next player's move can be surprisingly accurate (but still far from perfect) just by analysing local patterns and the next moves made within these local regions. In this sense the global state of the game had only a limited effect on the quality of choice for the next moves.

The choice of Go for this type of study, as opposed to another game such as Chess, is quite deliberate. Go has the very useful property of being highly spatially structured in the types of strategies employed by players. This is due to the fact that victory is decided on the basis of territory gained and that stones, once placed cannot be moved except when they are captured and that capture itself is a very spatial concept in Go. This lends Go very favourably to a type of localised spatial analysis that many other popular, games, particularly Chess with its in-homogeneous and moveable pieces, that are not susceptible. Popularity is also a critical ingredient, we need sufficient data in an electronic format in order to make these techniques work. With these considerations in mind, Go is a natural choice.

## Information Theory and Complex Games

When playing complex competitive games, such as board games or card games, one of the fundamental questions players ask themselves is: "What will the other player do next?" It is equally fundamental that what a player will do next is based on their level of skill. In this work we use Go to explore the uncertainty a player has regarding their opponent's next moves and how this changes as the players' level of skill increases. The tools of information theory are particularly well suited to our analysis as the principal unit is the "bit" providing a precise measure of how much information is available and how this information changes both within and between distributions.

Entropy is a well-established measure of the uncertainty an observer has in the outcome of a random event before the event is observed (Shannon 1948). Given a finite set of discrete events  $X = \{x_1, x_2, \dots, x_j\}$ , the outcome of a random event is labelled  $x$  and the probability of  $x_i \in X$  occurring is  $p(x = x_i)$  where  $\sum_{x_i \in X} p(x = x_i) = 1$ . We call the probability mass function  $p(x)$  and  $p(x = x_i)$  is written  $p(x_i)$ . The definition of entropy is then Shannon (1948):

$$H(p(x)) = - \sum_{x_i \in X} p(x_i) \log_2(p(x_i)) \quad (1)$$

where it is assumed  $0 \log(0) = 0$ . The minimum of (1) is zero and it is achieved when  $p(x_i) = 1$  for some  $i$  and zero otherwise. The maximum of (1) is achieved when  $p(x_i) = n^{-1}$  for all  $x_i \in X$  and  $n$  is the size of set  $X$ , i.e. for fixed  $n$  the entropy is maximised for the uniform distribution and minimised when the outcome is *deterministic*.

In order to motivate this definition for the game of Go, note that the entropy can be thought of as the expected value of  $-\log(p(x_i))$  over events  $x_i$ . The  $-\log(p(x_i))$  term is the amount of information that an observer gains when they see that  $x_i$  is the outcome of a random event. For example if  $p(x_i) = 1$  then  $-\log(p(x_i)) = 0$  as observing an outcome that is guaranteed to occur tells us nothing that is not already known before the observation was made. However if  $p(x_i) = 0.5$  then  $-\log_2(0.5) = 1$ , i.e. 1 bit<sup>2</sup> of information is gained having observed  $x_i$  as an outcome if  $x_i$  had a 50% chance of occurring. The entropy is then the average number of bits an observer expects to gain in observing a random event. This enables us to relate two important concepts: the uncertainty an observer has in the outcome of a random event before it occurs is equivalent to the expected amount of information they will gain after the event has occurred.

We consider two players; player one is about to make a move during a game of Go, player two is interested in where player one will choose to make that move. We define the strategic capacity of a population of players as the uncertainty player two has regarding player one's choice of next move if all that player two knows is that player one comes from a population with a known skill level. In the case where only the opponent's rank is known and nothing of their style or other idiosyncratic influences are known, this is the best estimate of the uncertainty player two has in player one's choice of next move. In this sense the population entropy aggregates all of the individual variation within the population into a single measure.

We now consider measuring the difference between distributions used by two different populations of players, the Kullback–Leibler (K–L) divergence (Kullback and Leibler 1951). To motivate the K–L divergence we consider two different probabilities for the one event  $x_a$  given by:  $p_1(x_a) = 0.25$  and  $p_2(x_a) = 0.5$ . The information gained having observed outcome  $x_a$  is:  $-\log_2(p_1(x_a)) = -\log_2(0.25) = 2$  bits and  $-\log_2(p_2(x_a)) = -\log_2(0.5) = 1$  bit. Given that entropy measures the expected information gain, a natural measure of the difference between two probabilities is to take the difference between the information gained having observed  $x_a$  using the different probabilities:

$$\begin{aligned} -\log(p_1(x_a)) + \log(p_2(x_a)) &= -\log(p_1(x_a)/p_2(x_a)) \\ &= 1 \text{ bit.} \end{aligned}$$

The expectation of the difference in information between each outcome over all outcomes is called the K–L divergence, given by the weighted difference for each event  $x_i$ , where the weighting is with respect to  $p_1(x)$ :<sup>3</sup>

<sup>2</sup> We will only use  $\log_2$  so the information measured is always in bits.

<sup>3</sup> Note that we make the usual assumptions of  $0 \log(0) = 0$  and  $0 \log(0/0) = 0$ .

$$K(p_1(x), p_2(x)) = \sum_{x_i \in X} p_1(x_i) \log \left[ \frac{p_1(x_i)}{p_2(x_i)} \right] \quad (2)$$

The larger the value of the K–L divergence the more significantly different one distribution is from another where zero bits difference only occurs if the distributions are the same for every  $x_i$ . As an example consider the distributions  $p_1(x_1) = 0.1$ ,  $p_1(x_2) = 0.9$  and  $p_2(x_1) = 0.2$ ,  $p_2(x_2) = 0.8$ . The K–L divergence is  $K(p_1(x), p_2(x)) = 0.0529$  bits. On the other hand the distributions  $p_1(x_1) = 0.9$ ,  $p_1(x_2) = 0.1$  and  $p_2(x_1) = 0.2$ ,  $p_2(x_2) = 0.8$  have  $K(p_1(x), p_2(x)) = 1.6529$  bits, i.e. the difference between the two cases is a factor of more than 30. This difference is a measure of the change in the *micro-structure* of the distributions in the sense that the entropy is constant for each example of  $p_1(x)$  and  $p_2(x)$  but the relative difference in which event is more likely,  $x_1$  or  $x_2$ , is significantly different in the second example. A more sophisticated version of this example is used in the Discussion to illustrate key aspects of our results.

Equation (2) is undefined when  $p_1(x_i) \neq 0$  and  $p_2(x_i) = 0$  for at least one  $i$  (Lin 1991). This can occur when one group of players makes a move another group never makes so instead we use the modified form given by Lin (1991):

$$\tilde{K}(p_1(x), p_2(x)) = \sum_{x_i \in X} p_1(x_i) \log \left[ \frac{2p_1(x_i)}{[p_1(x_i) + p_2(x_i)]} \right] \quad (3)$$

This can be seen as the expected difference in the information between  $p_1(x_i)$  and the *midpoint* of  $p_1(x_i)$  and  $p_2(x_i)$  for each  $i$  as measured in bits. Note that  $0 \leq \tilde{K}(p_1(x), p_2(x)) \leq 1$  where zero is only attained if  $p_1(x) = p_2(x)$  and 1 bit is attained if  $p_1(x_i) \neq 0$  when  $p_2(x_i) = 0$  for all  $x_i$ . We are interested in two particular calculations:  $\tilde{K}$  between players of successive skill levels,  $\tilde{K}(p_k(x), p_{k+1}(x))$ , and  $\tilde{K}$  between grand-masters and other skill levels,  $\tilde{K}(p_{gm}(x), p_k(x))$ . Roulston (1999) recently derived the error estimates used in this work.

## Methods

Building on the work of Stern et al. (2006) and Reitman (1976), we combine the notions of the cognitive “chunks” Reitman uncovered, as well as larger patterns of play studied in books on Go opening positions and strategies for corner regions, with the statistics for what moves are played next in the local context. Then we segregate players according to their known rank, as defined by the number of games they’ve won against other players of known ranks. This provides us with distributions over possible next moves from a known starting configuration. Comparing the way in which these distributions change as player ranks change provides us with a way of understanding both qualitative and quantitative information on how the collective strategies of populations of players change with levels of expertise.

<sup>3</sup> Note that we make the usual assumptions of  $0 \log(0) = 0$  and  $0 \log(0/0) = 0$ .

We derive our probability distributions from databases of game records, either from commercially available records<sup>4</sup> (games of professional players) or collected from online communities of Go players<sup>5</sup> (games of amateur players). From these sources we obtained game records for 160,000 players with skill levels from the threshold of expertise to the best in the world. Both the online and the professional games are played for rank position, however online games are of course different in the sense that they are not held in a formal setting with national and international reputations at stake. However, what we are able to show is that there are significant results observable in the transition points between amateurs and professional players. Amateurs may play differently in different settings but the densest accurate data sets obtainable for amateur players are those of online play.

The most junior ranked players had a rank of 2 kyu amateur and the most senior players had a rank of 9dan professional. An amateur rank of 2 kyu can often be achieved with a year or two of recreational play whereas the rank of 9dan professional will take decades of dedicated training. In increasing level of skill, the set of ranks is: 2 kyu amateur, 1 kyu amateur, 1dan amateur, 2dan amateur, . . . , 8dan amateur, 1dan professional, 2dan professional, . . . , 9dan professional. We abbreviate these names as: A2K, A1K, A1Dan, A2Dan, . . . , A8Dan, P1Dan, . . . , P9Dan. In order to collect sufficient data for each level of skill, we aggregated some of these ranks so the list of skill levels is: A2K, A1K, A1-2Dan, A3-8Dan, P1-4Dan, P5-7Dan and P8-9Dan.


Each of these collections of games was then used to search for commonly occurring, well known and well studied patterns of play from which standard moves are often made. There are three broad categories to these patterns. The smallest and most local starting patterns are conceptually similar to the chunks described by De Groot and Gobet (1996) and based on the work of Reitman as discussed above. These patterns can occur anywhere on the board at almost any point in the game and so are not spatially or temporally localised. See Fig. 1 for the patterns used. The first two stones are unnumbered, they are the starting patterns, the stones are then numbered according to the sequence in which they were played. Note that other moves may have been made elsewhere on the board between each (local) placement of stones.

The sequence of numbered stones from 1 through 5 in Fig. 1 are the most probable choice of moves made by the 8-9Dan professionals. The numbers appearing on the vacant vertices are all of the alternative options for the next move (in this case the sixth move) recorded in the games database for the 8-9Dan professionals. Move option 1 is the most frequently occurring next move, option 2 the second most frequent etc. The actual probabilities associated with these possible next moves make up the distributions used to calculate the information measures. There are *always* fewer next moves than there are empty vertices in the local patterns used and the probability of the next move being in one of these positions is considerably different from a uniform distribution. The same type of analysis has been applied to the other figures discussed next.


The mid-sized patterns are spatially localised around the corners and typically occur in the early stages of the game, see Fig. 2. Importantly, the smallest patterns

<sup>4</sup> From the GoGod collection: <http://www.gogod.co.uk>.

<sup>5</sup> From the KGS Go server: <http://www.gokgs.com>.



**Fig. 1** The two smallest patterns used. The starting pattern (the *two black un-numbered stones*) on the *left* is called ‘skip one’ and the starting pattern on the *right* is called ‘knight’s move’ in Reitman’s study



**Fig. 2** The two mid-sized patterns used. The pattern on the *left* is a variation of the ‘avalanche’ joseki and the pattern on the *right* is a variation of the ‘4–4 point low approach high extension’ joseki. Note that the board is bounded by the corner in these patterns so that they only ever occur in the corner of the board, unlike the smaller patterns in Fig. 1 that may occur anywhere on the board, including the corners

described above are often used to define the starting positions within these larger corner regions. Each mid-sized pattern is a *Joseki* in the terminology of Go (literally: “fixed stones” or “set pattern”) and there are many well studied variations of play that occupy novices and masters throughout their playing lives. The choice of which variation to play is at least partly contingent upon the context of the rest of the board.

The largest patterns, called *Fuseki*, encompass the whole board and frequently occurring positions only occur within the very early part of the game. Note though that the smaller corner positions often define local patterns of play within this whole board context. In this sense each smaller pattern plays a role in the larger patterns, providing a conceptual embedding of the patterns within each other.

## Results

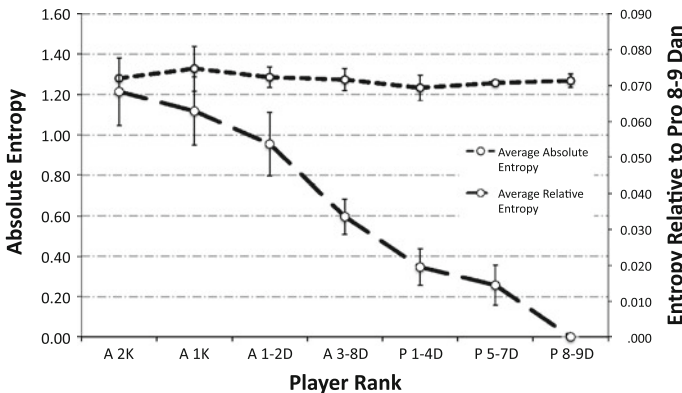
Figure 3 shows the strategic capacity and  $\tilde{K}$  averaged over all positions considered. The strategic capacity is flat all the way from pre-expert to grand-master. This is



surprising. One might think that an increase in player skill would correspond to a decrease in move predictability. In other words, the information gained from learning the move of a master would be greater than for a beginner. However, for basic positions, this is not the case. In any given position, there will be many empty positions where a move would never be made.

For 256 empty positions (maximum is 361 at the beginning of the game) this would correspond to 8 bits of information to define where a move was made if all positions are equally likely, or about 5 bits for the smaller board sections. Thus, the information gained from knowing where a move is made is much less than would be expected from a uniform distribution across all empty positions. Figure 3 shows this information from the actual game data, where the information is found to be just 1.3 bits. But, remarkably, this is constant across all ranks and moves. So the capacity of the strategy space remains the same but the organisation within it changes. In the Discussion a simple way to satisfy such a constraint is considered.

Figure 3 also shows the information divergence, again averaged over all positions. It decreases monotonically as expected but is almost linear. Figure 4 shows the same data, but now split into two categories: Fuseki, the moves made in the early part of the game and involve the whole board; and the smaller patterns where local fighting often occurs. Firstly note that  $\tilde{K}$  of the smaller patterns decrease as before but the slope is smaller. The Fuseki on the other hand remain flat out until expert level (A1-2Dan) then fall steadily to grand-master. The implication is that the local tactical fighting is learned early and improves, albeit slowly, with increasing skill. However, for the global Fuseki there is a pronounced lack of *any* progress until expert. Thus there is a transition at expert level, where local skill is



**Fig. 3** Strategic capacity and information divergence as a function of player rank. For the grand-master average relative entropy curves, we took the 34 different board configurations and for each rank we generated probability distributions of what next move was played. These 34 configurations included very local regions (3 × 5 regions) up to the whole board (19 × 19). They also included continuations (up to 6 move sequences) from a given starting position. These 34 configurations, and their associated probability distributions, allows for a very broad sweep of many different types of well known strategic positions. Each of these probability distributions was then compared, using (3), with the grand-master probability distributions of what next move was played. The result is the Average Strategic Divergence. Errors are ± s.e.m

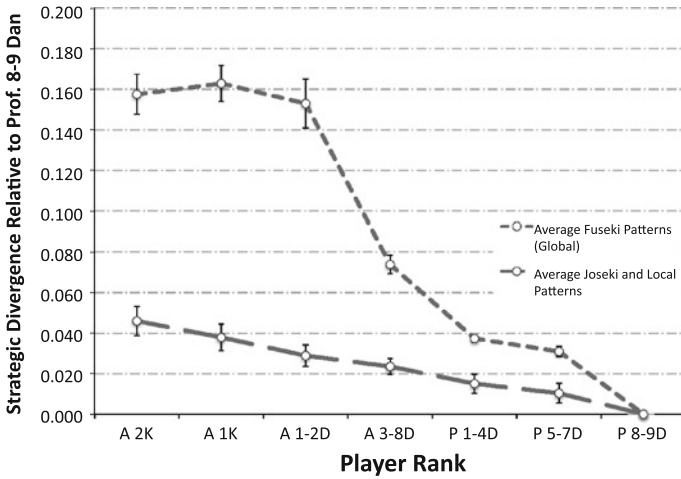


Fig. 4 The strategic divergence for global and local patterns. Errors are  $\pm$  s.e.m

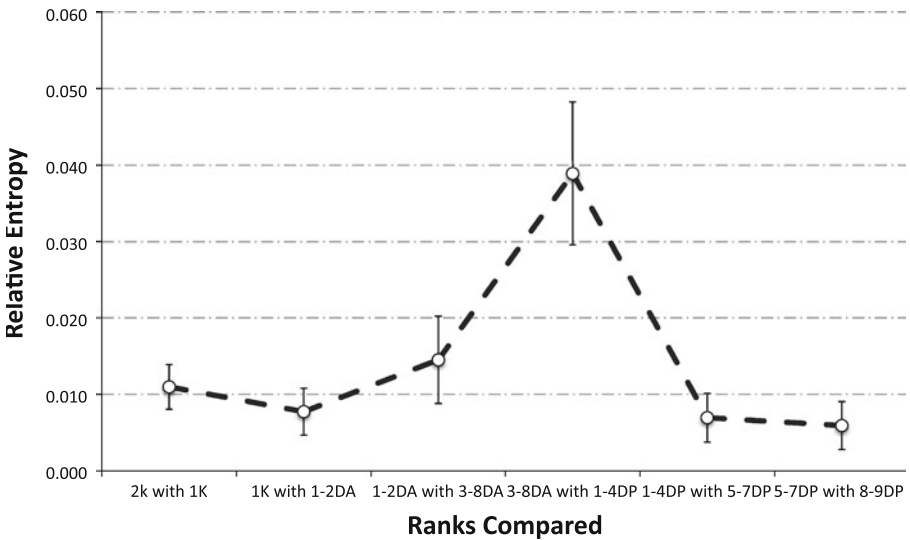


Fig. 5 The incremental strategic divergence. Details as in Fig. 3. Errors are  $\pm$  s.e.m

sufficiently developed such that awareness of global strategies is sufficiently developed for progress to be made towards grand-master like choices.

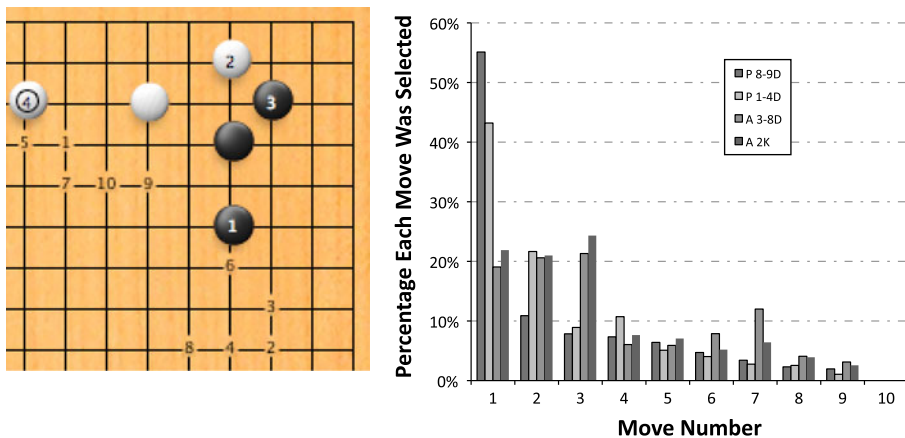
Figure 5 shows the incremental divergence and reveals a second surprising transition. At the transition to professional level there is a pronounced peak. Taken in the context of Fig. 3 this means that there is a significant change between ranks at this point.

**Discussion**

The results presented in this paper are the first to measure the development of human expertise for an advanced cognitive domain over a large cohort of people who have been practising for decades. The two transitions observed are typical of the way expertise in general is thought to develop as reported by Ericsson and Lehmann (1996). We speculate that these transitions reflect a fundamental shift in the way in which players understand and play the game as their level of expertise increases. While Ericsson and others have previously studied on the qualitative nature of these transitions, this novel approach, applied to entire populations of players, is able to show their occurrence quantitatively for the first time.

In order to understand the nature of the second transition point, we plot the next-move distribution for one particular board configuration for four different ranks of players: the lowest ranks (A2K), the two ranks at which, when compared, the transition occurs (A3-8Dan and P1-4Dan) and the grand-master rank (P8-9Dan), see Fig. 6. These distributions are the next-move probabilities in a  $7 \times 7$  region of the board in the corner, where there are only two stones already played in this region. This initial configuration was chosen as it is very well known and quite thoroughly studied by Go players.

The first observation is that the first three moves for A2K and A3-8Dan have very similar probabilities, approximately 20%. Consequently these three moves would contribute very little to the divergence between these two ranks as calculated by (3). It can also be seen that the second choice of move for P1-4Dan is also very similar to that of the two lower ranked players. However, the first move choice for P1-4Dan is significantly different from that of the two lower ranked players. This single move



**Fig. 6** The board state and probability distributions over the next moves. *Top* One of the joseki showing the first six stones played in the local area by 8-9Dan professionals. *Bottom* Four example histograms of the frequency each of the ten moves might follow. Note the order of the moves on the horizontal axis in this plot is with respect to the 8-9Dan professionals, i.e. the most preferred move, labelled ‘1’, is most preferred by the 8-9Dan professionals, the move labelled ‘2’ is the second most preferred by these players etc

distinction contributes almost two thirds (65.1%) of the total incremental divergence between P1-4Dan and A3-8Dan. When comparing P1-4Dan with P8-9Dan it can be seen that there are particular differences in move distribution. However, the progression from P1-4Dan to P8-9Dan is a series of much smaller and more incremental changes in the player's move distributions.

Learning this simple difference between the choice of where to place the third stone in a very open region of the board expresses the amount of knowledge a player accumulates when moving from the best amateur ranks to the junior professional ranks. But the choices made by the players cannot simply be based on the information available in the local region: there are only two other stones, it is highly improbable that players who have been observing the same two stone pattern for many years suddenly gain some great insight into these two stones that had previously eluded them. Instead it is almost certainly true that it is insights into the broader context of the game that enables the players to suddenly draw such a sharp distinction between one choice of move over another.

It is possible to replicate the key findings of this work for the amateur-professional transition with four different distributions representing four levels of player's skill where each distribution is composed of the same three choices. The qualitative features we require of these distributions are:

1. The entropies are all approximately the same, emulating the constant strategic capacity of the populations.
2. The divergences between the most skilled player's distribution and the other three be decreasing slightly in value as the skill of the groups increases representing the smoothly decreasing divergence between increasingly skilled players and the most skilled players.
3. The divergence between the two least skilled groups be close, emulating the similarity between groups of amateur players.
4. The divergence between the two middle ranked groups are significantly different from the two lowest ranked groups, emulating the sharp transition from amateur to professional.

Consider the set of possible choices  $X = \{x_1, x_2, x_3\}$  and four different populations (groups) of players denoted  $G_i$  where the relative skill levels of each group are:  $\text{skill}(G_1) < \text{skill}(G_2) < \text{skill}(G_3) \ll \text{skill}(G_4)$ , i.e. the most skilled group is considerably more skilled than the other three groups,  $G_1$  and  $G_2$  represent amateurs,  $G_3$  and  $G_4$  represent professionals and  $G_1, G_2$  and  $G_3$  are immediately adjacent in skill levels. Each group has associated with it a probability mass function over  $X$ :

$$p_1(x) = \{0.5, 0.00005, 0.49995\}$$

$$p_2(x) = \{0.5, 0.005, 0.495\}$$

$$p_3(x) = \{0.5, 0.49, 0.01\}$$

$$p_4(x) = \{0.77, 0.115, 0.115\}.$$

Using (1) the entropies of these distributions (the strategic capacities of the groups) are:  $H(p_1(x)) = 1.0007$  bits,  $H(p_2(x)) = 1.0404$  bits,  $H(p_3(x)) = 1.0707$  bits and  $H(p_4(x)) = 1.0080$  bits. These entropies are approximately equal despite the

quite different composition of the individual event probabilities. Using (3),  $\tilde{K}$  between these distributions are:

$$\tilde{K}(p_4(x), p_1(x)) = 0.1659$$

$$\tilde{K}(p_4(x), p_2(x)) = 0.1603$$

$$\tilde{K}(p_4(x), p_3(x)) = 0.1548$$

$$\tilde{K}(p_1(x), p_2(x)) = 0.0033$$

$$\tilde{K}(p_2(x), p_3(x)) = 0.4526$$

It is readily seen that the  $\tilde{K}$  of  $p_2(x)$  with  $p_3(x)$  gives a very large value when compared with that of  $p_1(x)$  with  $p_2(x)$ , emulating the developmental shift observed in going from amateur play to professional play. On the other hand all three lower ranks have similar  $\tilde{K}$  that decreases slightly with increasing skill when measured relative to the highest ranked group, emulating the small incremental decrease observed when comparing lower ranks to the grand masters in our data. These artificial distributions suggest what may be occurring in our original data; a significant restructuring of which move is preferred over another. In this example the entropy of  $G_3$  is similar to that of  $G_2$ , but  $G_3$  prefers  $x_2$  much more than  $x_3$ , a reversal of preferences compared to  $G_2$ . When looking at the original data it is precisely this sort of transition in the composition of the probabilities that result in the transition in going from amateur to professional: a significant reorganisation in the micro-structure of the move preferences between *adjacent* skill levels that is not apparent by simply observing the entropies *or* by comparing the low skilled distributions with those of the grand-masters.

In light of these findings we consider what it means to have a probability distribution over next moves. It might be thought that the population's distribution accurately reflects some internal uncertainty each player within that population has regarding where they might move next, and that if a single player were given the same local pattern repeatedly the same distributions observed in this study might emerge over the next moves. This seems unlikely given that players learn a range of possible alternatives for any given local pattern but they choose one particular strategy (sequence of moves) given what they understand of the game as a whole, not based simply on the smaller local patterns explicitly used in this study (except for the whole board openings).

Consequently the probability distributions reflect the influence the rest of the board plays on the local decision of which move should be chosen next. In this view, for any simple, small pattern of stones it is possible that good players, such as those considered in this study, could have learned the optimal (purely local) move to make. This would imply that the local patterns were independent of the global context in which they occurred. Instead we observe a relatively constant degree of uncertainty in which move should be made next across all skill levels, implying that the uncertainty produced by the global context does not change despite the many different actual but unobserved contexts in which the moves occurred (except for the whole board openings).

Furthermore, this uncertainty is identical with how much information the rest of the board conveys regarding the local decision. If only one move were ever made from a given local pattern then we might reasonably hypothesise that the local pattern provides all of the information necessary to disambiguate all of the moves and so make an optimal choice. The entropy of the distribution of next moves would then be zero and therefore zero information is required from the rest of the board. On the other hand if all possible next moves were equally likely so that the entropy is maximised then we could infer that the local pattern conveys *no* information that is useful to disambiguate where the next move should be made and only the global context in which the pattern occurs, and nothing about the pattern itself, is relevant. This also emphasises what we observe in the second transition: a sharp discontinuity in the way in which the global context is able to differentiate between local move choices.

After averaging over tens of thousands of games we might expect that all move choices would be equally likely (i.e. maximum entropy). For example, the two smallest patterns may occur anywhere on the board and at almost any point in the game. However, across all ranks of players and every sequence of move distributions, the first two most frequently played moves accounts for  $68.5\% \pm 2.75\%$  (mean  $\pm$  s.e.m) of all choices observed. This provides strong support for Reitman's (1976) findings regarding these patterns being strategically useful. It also suggests that these patterns might connect with the broader context of the game in two commonly occurring ways.

We suggest that the results observed here for the distinct, sharp changes in player behaviour are equivalent to the phase transitions observed in many developmental studies. Ericsson and colleagues have made extensive studies of expertise and its acquisition in fields as diverse as musicians and athletes (see Ericsson and Lehmann (1996) for a review). Of principal note is the work on the 'Three Phases of Development of Expert Performance' [see Ericsson and Charness (1994, Figure 3, p. 739)]. Here Ericsson and Charness describe these three phases as: (1) initiation, (2) transition to full time involvement and (3) experts seeking to make eminent achievements. These transition points between phases are qualitative descriptions of the distinct quantitative transitions seen in our results: distinct transition points in the development of expertise based on the strategies employed by the players themselves.

The consequences of these results are also not trivial when considering artificial systems such as supervised learning with, for example, neural networks. The most senior players in our database might be thought of as suitable 'oracle' for such learning, the best players in the world and therefore ideal candidates whose strategy a neural network might want to approximate as it learns. As Fig. 3 shows, lower ranked players can apparently approach the oracle's distribution of move choice smoothly. This is deceptive, however, as shown in Fig. 5. The amount of information required to transform one adjacent distribution into another is not guaranteed to be as smooth as that of the changes observed relative an oracle might imply. Curiously these results also imply that better experts *do not* necessarily use more information from the global context, so more (global) information is not *necessarily* required in order to be a better player, although it might be very useful,

instead what information is attended to and how the global context is configured is what changes as expertise is acquired. While these findings are not necessarily novel in the psychological literature or even to the AI community, such a clear quantitative demonstration that is able to inform both fields is highly novel.

Finally, we ask whether results obtained for a complex game such as Go are applicable to the execution of expertise in general. There is already an established tradition of studying expertise through Chess and Go is its natural successor (Burmeister and Wiles 1995). In countries such as Japan and Korea where Go is of major cultural significance, children learn to play early and play intensely. Their learning of the game is of a similar time frame, time scale and intensity to learning a language. We conjecture that the few games which are successful over millennia match human cognitive aptitudes and development especially well. The methodology we describe here can be used to search for knowledge transitions in any domain where a large volume of data is digitally available. Other online games such as poker are clearly candidates for future work, but potentially professional areas such as medicine or financial trading could benefit from the method of analysis used in this study.

## References

- Atherton, M., Zhuang, J., Bart, W., Hu, X., & He, S. (2003). A functional MRI study of high-level cognition. I. The game of chess. *Cognitive Brain Research* 16(1), 26–31.
- Brooks, R. A. (1990). Elephants don't play chess. *Robotics and Autonomous Systems* 6(1–2), 3–15, designing Autonomous Agents.
- Burmeister, J., & Wiles, J. (1995). The challenge of Go as a domain for AI research: A comparison between Go and chess. In: *Proceedings of the third Australian and New Zealand conference on intelligent information systems*.
- Buschkuhl, M., Jaeggi, S., Hutchison, S., Perrig-Chiello, P., Däpp, C., Müller, M., et al. (2008). Impact of working memory training on memory performance in old-old adults. *Psychology and Aging* 23(4), 743–753.
- Campbell, M., Hoane, A., Jr., & Hsu, F. (2002). Deep blue. *Artificial Intelligence*, 134(1–2), 57–83.
- Chase, W., & Simon, H. (1973). Perception in chess. *Cognitive Psychology* 4(1), 55–81.
- Chen, X., Zhang, D., Zhang, X., Li, Z., Meng, X., He, S., et al. (2003). A functional MRI study of high-level cognition: II. The game of GO. *Cognitive Brain Research* 16(1), 32–37.
- Deppe, M., Schwindt, W., Kugel, H., Plassmann, H., & Kenning, P. (2005). Nonlinear responses within the medial prefrontal cortex reveal when specific implicit information influences economic decision making. *Journal of Neuroimaging*, 15(2), 171–182.
- Dreyfus, H. (1992). *What computers still can't do: A critique of artificial reason*. Cambridge: MIT Press.
- Dreyfus, S. (2007). Why expert systems do not exhibit expertise. *IEEE Expert* 1(2), 86–90.
- Ekbia, H. (2008). *Artificial dreams: The quest for non-biological intelligence*. Cambridge: Cambridge University Press.
- Engle, R., Tuholski, S., Laughlin, J., & Conway, A. (1999). Working memory, short-term memory, and general fluid intelligence: A latent-variable approach. *Journal of Experimental Psychology: General*, 128, 309–331.
- Ericsson, K., & Charness, N. (1994). Expert performance: Its structure and acquisition. *American Psychologist* 49, 725–747.
- Ericsson, K., & Lehmann, A. (1996). Expert and exceptional performance: Evidence of maximal adaptation to task constraints. *Annual Review of Psychology* 47, 273–305.
- Ericsson, K., Krampe, R., & Tesch-Römer, C. (1993). The role of deliberate practice in the acquisition of expert performance. *Psychological Review—New York* 100, 363–363.

- Gamez, D. (2008). Progress in machine consciousness. *Consciousness and Cognition* 17(3), 887–910.
- Gobet, F. (1998). Expert memory: A comparison of four theories. *Cognition* 66(2), 115–152.
- Gobet, F. (2005). Chunking models of expertise: Implications for education. *Applied Cognitive Psychology* 19, 183–204.
- Gobet, F., & Simon, H. (1996). Recall of rapidly presented random chess positions is a function of skill. *Psychonomic Bulletin & Review* 3(2), 159–163.
- De Groot, A., & Gobet, F. (1996). *Perception and memory in chess*. Assen: Van Gorcum and Co.
- Jaeggi, S., Buschkuhl, M., Jonides, J., & Perrig, W. (2008). Improving fluid intelligence with training on working memory. *Proceedings of the National Academy of Sciences* 105(19), 6829.
- Kullback, S., & Leibler, R. (1951). On information and sufficiency. *The Annals of Mathematical Statistics*, 22(1), 79–86.
- Kuo, B., Yeh, Y., Chen, D., Liang, K., & Chen, J. (2008). The capacity constraint in the prefrontal and parietal regions for coordinating dual arithmetic tasks. *Brain Research*, 1199, 100–110.
- Lin, J. (1991). Divergence measures based on the Shannon entropy. *IEEE Transactions on Information theory* 37(1), 145–151.
- Newman, L., & Polk, T. (2008). The computational cognitive neuroscience of learning and memory: Principles and models. *Advances in Psychology*, 139, 77–99.
- Reitman, J. (1976). Skilled perception in Go: Deducing memory structures from inter-response times. *Cognitive Psychology* 8(3), 336–356.
- Ross, P. (2006). The expert mind. *Scientific American*, 295(2), 46–53.
- Roulston, M. (1999). Estimating the errors on measured entropy and mutual information. *Physica D: Nonlinear Phenomena* 125(3–4), 285–294.
- Shannon, C. (1948). A mathematical theory of communication. *Bell System Technical Journal* 27, 379–423.
- Simon, H. (1973). The structure of ill structured problems. *Artificial intelligence* 4(3–4), 181–201.
- Simon, H. (1991). Bounded rationality and organizational learning. *Organization Science* 2(1), 125–134.
- Stern, D., Herbrich, R., & Graepel, T. (2006). Bayesian pattern ranking for move prediction in the game of go. In: *Proceedings of the 23rd international conference on machine learning* (Vol. 148, pp 873–880).
- Tromp, J., & Farnebäck, G. (2006). Combinatorics of Go. In: *Proceedings of the 5th international conference on computers and games* (pp 84–99). Berlin: Springer.
- Wang, P. (2007). Three fundamental misconceptions of artificial intelligence. *Journal of Experimental and Theoretical Artificial Intelligence* 19, 249–268.
- Wang, Y., & Gelly, S. (2007). Modifications of UCT and sequence-like simulations for Monte-Carlo Go. In: *IEEE symposium on computational intelligence and games* (pp. 175–182). Honolulu, Hawaii.