Global Context Influences Local Decisions

Terry Bossomaier
CRiCS (Centre for Research in Complex Systems)
Charles Sturt University
Australia.
Email: tbossomaier@csu.edu.au

Michael Harré
CRiCS, Charles Sturt University
& Centre for the Mind, University of Sydney
Australia.
Email: mike.harre@gmail.com

***Abstract*—This paper studies the development of human expertise in the game of Go. Although superficially a simple game, Go is the most difficult of all established games for artificial intelligence, no computer program yet achieving top international level on a full 19x19 board. On smaller boards, such as 9x9 computers *are* competitive, implying that the understanding of the complex global interactions is the key to human superiority.**

By mining thousands of positions online, we show that at some player levels the sequence of plays leading up to a *local* position is a stronger determinant of the next move than the position alone. This suggests that the sequence of plays is an indicator of global strategic factors and thus provides a context for the next move in addition to the local position itself.

***Keywords*-game of go; decision making; entropy; online data mining**

I. Introduction

The big picture often influence or override local factors in many areas of human expertise, from board games to politics. Challenging games, such as Chess and Go, provide an excellent framework for studying expertise [1], [2], [3], [4] since they are both strategically deep but tightly constrained. This paper presents a striking demonstration of this, mined from thousands of decisions online. In recent work we have demonstrated transitions in the acquisition of expertise in the game of Go [5]. This game is interesting because it is currently the most difficult of all established games for computational intelligence. Unlike Chess, where the IBM computer Deep Blue [6], [7] triumphed over world champion Kasparov.

We also demonstrated therein, from calculation of mutual information between moves, that one of these transitions has the character of a *phase transition* [8]. The idea of a phase transition comes originally physics, such as the melting of ice to give water. When such a physical phase transition occurs there is a dramatic reorganisation of the system. In this case water molecules which were fixed rigidly in place in ice become free to move around and perhaps travel long distances. During a phase transition, systems exhibit long range order, where there are correlations in activity or structure over large distances and system parameters often exhibit power law behaviour, or fat-tailed distributions. Another example of a phase transition is in the formation of random graphs. At the transition the average path length, in other words the number of steps from one node in the graph to another rises to a peak, and then drops back down again.

A dynamical system examples is the Vicsek model developed for studying magnetic transitions in solid state physics [9]. In this model particles travel around a two dimensional grid, and, when they come within some specified distance of each other, their directions of movement get slightly closer together. Phase transitions occur in this system as particles flow around in groups, like flocks of birds, but these groups are dynamic, continually forming and dissolving.

Mutual information is a system property which measures the extent to which the structure or behaviour of one part of a system predicts the behaviour of another. In the Vicsek model of above, at the transition the direction and velocity of one particle provides some information about the direction of all the other particles. The mutual information peaks during the phase transition [9], [10], and this is thought to be a general property of phase transitions along with the other characteristics, notably long range order and power law characteristics. We found a peak in mutual information as a function of rank amongst Go players from 1 Dan Amateur through to the very top players, 9 Dan Professional [8]. The previous work [8] has demonstrated phase transitions in collective human decisions in Go. This paper presents evidence that there is global influence on local decisions and that the influence is greatest during the phase transition.

A. State of the Art in Game Expertise

Much of the work on human expertise has been based on games, especially Chess, as in Gobet's extensive work [1], [11]. One of the key ideas, essentially from Nobel Laureate Herbert Simon, is that human expertise involves building a huge library of patterns [12], [13], although the application of these ideas in artificial intelligence for games is relatively new [14].

These patterns build up through the formation of *chunks*, and psychological observables, such as memory for Chess positions are well predicted by models such as CHREST [3]. The way the cognitive structures in the brain might change as expertise develops, and in particular the appearance of phase transitions, is relatively new introduced by Harré and Bossomaier [8], [5].

Further recent advances have been limited, particularly in Go where a combination of the game space complexity of Go [15] and a lack of genuinely human like heuristics such as an evaluation function make progress difficult. However with the development of ever more effective random sampling techniques, such as the UCT-Monte Carlo approach currently favoured by AI system developers [16], some progress has been made in achieving strong amateur play. However these techniques do not address the inherent complexity of the game and the techniques that humans have developed in order to address such issues, almost completely because such techniques are not subject to easy investigation.

We argue that the sources of information players use in order to make good decisions are of two types: *local* and *global.* While every level of player in our study has learned a great deal about the game of Go over the course of their lives, it is how this information is implemented via the choices they make that is of interest to us. This relevance of the division of the problem space in to these two parts can be seen in the work of Stern et al [17]. They were able to produce 'best in class' move prediction for professional players in Go, achieving a 34% success rate. This was achieved by training their system on 181,000 expert game records and using the most modern techniques available for their analysis. The level of success achieved in this work highlights one of the principal difficulties of good performance in complex tasks: exact pattern matching is not enough; AI systems need to be able to model how non-local aspects i.e. information that cannot be derived by exactly matching board configurations, influence decisions. Loosely interpreted this is what is called *influence* in Go and before our recent work it had not been reported in the research literature.

II. Methods

Harré and Bossomaier [8] examined the game trees 6 moves (i.e. 3 black and three white) deep for around 8,000 games across a range of Go expertise. At the low end were 2 kyu Amateurs, a rank reached by serious players after a couple of years club play through the highest amateur rank of 6 Dan Amateur (6A) to the top professional rank of 9 Dan Professions (9P). Game data was obtained from the *pandanet* Go server. Full details of experimental procedures are given in Harré et al. [8]. The game trees were computed from 7x7 board sections in the corner, from games played between players of the same rank. No symmetry was exploited, apart from rotations to align each of the four corners (used to maximise data yield per game). Note that although these are the first 6 moves played in the region, they are not necessarily, and usually are not, the first 6 moves of the game.

For each possible move, m_i, three probability distributions were computed

1) the probability of the move, m_i, occurring, $P(m)$
2) the conditional probability, $P(m_I|q_i)$, of the move, m_i occurring from a given position, $q - I$
3) the conditional probability, $P(m_I|s_i)$ of the move occurring from a given position, *reached by a particular order of moves,* s_i

From these results the entropy and mutual information were calculated, but this paper addresses findings from the entropies alone. The move entropy, $H(M)$, is taken over all moves which can arise at each level in the game tree (i.e. for the 6 moves in the sequence). (eqn. 1).

$$H(M) = -\sum_i p(m_i)\log_2[p(m_i)] \tag{1}$$

Entropy is a measure of disorder or randomness and is maximal when the probabilities or all events are the same. When the first move in the region there are 49 possible positions and after 5 moves, 44, giving a maximal entropy of $\log_2 44 = 5.5$ bits. But since the moves are far from random the measured entropies are much lower than this.

The conditional entropy, $C(M|q_i)$, is the move entropy calculated from the moves which can arise in a given context, such as a position, q_j, or sequence of moves, s_j, leading to a position.

$$C(M|Q) = -\sum_i p(m_i|q_j)\log[p(m_i|q_j)] \tag{2}$$

with the same expression used for an ordered sequence of moves with s_j replacing q_j.

These entropic quantities are now calculated across all ranks from amateur 2q, denoted *am2q* through, the amateur dan ranks to *am6d*, and then to the highest rank of all, professional 9 dan, *pr9d* and are shown in figures 1–3.

III. Results

Figure 1 summarises the key findings of the paper. It shows the conditional entropy as a function of move in the sequence of 6 averaged across all ranks, both amateur and professional. Error bars are calculated as in Harré et al. [5]. Up to move 3 the entropy for both the ordered and unordered cases are the same. At move three they fall dramatically, but the ordered average falls about a third more.

Figure 2 shows the entropy at each move from a given position. For purely random moves the entropy at each move in the sequence would be between 5 and 6 bits. The entropies observed are, of course, much lower, usually less than 2 bits, reflecting the structure inherent in the game.

The entropy for the third move is slightly less than for the first two, but the entropy falls a little for the fourth move and a lot more for the fifth and sixth moves. This is not surprising given the reduced options available as the number

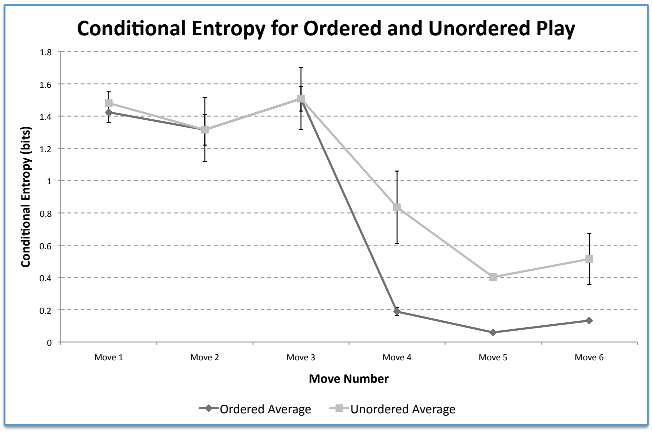

Figure 1. Conditional entropy (eqn. 2) as a function of move averaged over all ranks

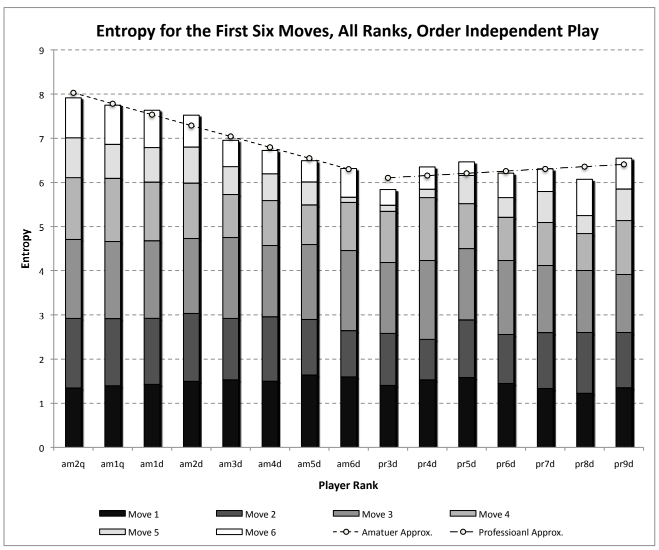

Figure 2. Entropies for moves from a given board position (reprinted from [8])

of stones on the board increases. The entropy summed over all moves declines linearly to the maximum amateur rank and then increases slightly from from the first professional rank.

Figure 3 shows the entropies which result from positions which arose from a particular sequence of play. These entropies are around 3 bits smaller, than for the unordered case. The slope of the regression line for the amateur levels is not so large, but the trend for the professionals displays a different pattern: the summed entropy jumps near the start of the professional ranks and then *decreases* with rank up to 9P with a slope very similar to that for the amateur ranks on the left of the figure.

The most interesting thing about this figure, though, is the way the entropy for the last three moves shrinks and vanishes as the amateur rank increases from 4A to 6A. In

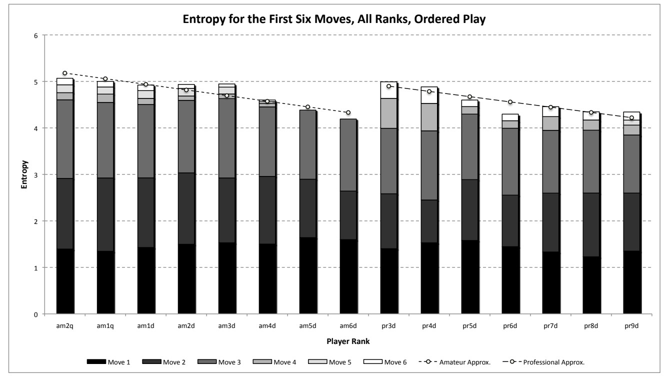

Figure 3. Entropy for the first six moves shown as a stacked bar chart. The black bars represent the entropy at move 1, the dark grey at move 2 ad so on for all six moves. The dashed regression lines show the total entropy for the amateur and professional sequences.

fact the summed entropy for the first three moves is quite similar to the unordered case, so the three bit loss is almost all in the last three moves.

IV. Discussion

There are two very interesting features of these results, which we consider in turn: the difference between ordered and unordered play; and the way the conditional entropy varies with rank.

That the ordered and unordered play differ, implies that the position at each move is *not* the sole determinant of the opponent response. The much lower conditional entropy after the first three moves for the ordered case strongly suggests that the sequence of moves has revealed something of the global context which has in turn fed back into move selection. To see this, imagine that black is strong in one area of the board and white in another. Since communication lines are of great strategic importance in Go, the locations of these areas will strongly influence the order of moves made in the local area we examine. The first three moves implicitly contain some of this information, which subsequently reduces the range of options in the second three moves.

The gradual decline in entropy with rank for amateur and professional reflects a gradual reduction in the space of range of options, which we could see as the elimination of poor moves in established situations, similar to the mastering of the opening in Chess.

Our data and results are explicitly based on an analysis of the local information, but by implication they also say a great deal about the global context that influences these localised decisions. The first three moves in our study have a reasonably similar conditional entropy of about 1.4-1.6 bits of information. This is the amount of information that is common between each successive move within the local region. Such measures of information are the best estimate of how much one stochastic variable can tell us about another [18]. The only other source of information available to the players are the pieces on the board that were not included within our local region. We exclude the possibility of being able to read the other opponent. While it is a debated issue as to the importance of such skills, we believe that it is much less significant than all the other pieces on the board that were not within the local area of study. The changing influence that non-local information has on decisions during a game, is evident in the the significant drop-off in the conditional entropy after move 3 in Figure 1, a drop-off of shared information from one move to the next of nearly an order of magnitude for the ordered play and about half that for unordered play.

This change in the levels of conditional entropy as the game progresses in the corner region of the board might be due to the reducing size of the move space as the board fills up. Will this might have some minor influence on our results, we should also expect such changes to be almost linear as the number of available positions only drops by a total of 1/49 per move. It is also possible, but exceptionally unlikely, that after move 3 players choose much more randomly, i.e. without concern for the pieces on board either local or non-local, than they did for the first three moves. Considering the vast training literature available to players that readily teach them the many different variations of the first 6 moves within the corner, and then how to contextualise these decisions by considering what pieces occupy nearby areas, we consider this to be an unlikely strategy.

Instead we argue that it is just this external influence, the influence of the stones arrayed on the rest of the board that is having such a striking influence on the condition entropy. This is perhaps not so surprising when considered in the light of the state of the game itself after 3 moves have been played in the corner. These first moves can be thought of as establishing the game board layout in terms of an 'opening book', highly stylised moves of local pieces where the local pattern can be thought of as essentially uncoupled from the rest of the board, or at least coupled to the same extent for these first moves. This coupling then changes significantly from the 4th move onwards where greater consideration needs to be afforded to the other pieces on the board. This change in focus of the information effectively reduces significantly the information coupling between the local moves and the local stones on the board.

The complete disappearance of entropy at the high amateur ranks is very interesting. It suggests that the at this level play has become somewhat stereotyped, and a major change in thinking is needed to advance, which indeed seems to happen on turning professional. Thus this loss of entropy is consistent with the long range order found in phase transitions [8]. This accords with the findings in Harré et al [8], wherein a peak in mutual information is found at the transition to professional, indicating some sort of major cognitive reorganisation. At present we do not know how to quantify such a reorganisation and this remains an exciting open question. Ongoing work is attempting to apply the CHREST models to Go [3] and to determine how phase transitions might be predicted.

ACKNOWLEDGEMENTS

This work was supported by the Australian Research Council under Discovery Project DP0881829 and the US Airforce under grant 104116.

REFERENCES

[1] F. Gobet, A. d. Voogt, and J. Retschitzki, *Moves in Mind*. Psychology Press, 2004.

[2] A. D. Groot and F. Gobet, *Perception and memory in chess: Heuristics of the professional eye*. Assen: Van Gorcum, 1996.

[3] F. Gobet, P. Lane, S. Croker, P. Cheng, G. Jones, I. Oliver, and J. Pine, "Chunking mechanisms in human learning," *Trends in Cognitive Sciences*, vol. 5, pp. 236–243, 2001.

[4] K. Ericsson and N. Charness, "Expert performance: its structure and acquisition," *American Psychologist*, vol. 49, pp. 725–7247, 1994.

[5] M. Harré, T. Bossomaier, C. Ranqing, and A. Snyder, "The development of expertise in a complex environment," *Minds and Machines*, vol. in press, 2011.

[6] X. Cai and D. Wunsch, "Computer Go: A grand challenge to AI," in *Challenges for Computational Iintelligence*. Springer Berlin, 2007, pp. 443–465.

[7] M. Campbell, A. Hoane, and F. Hsu, "Deep blue," *Artificial Intelligence*, vol. 134, no. 1-2, pp. 57–83, 2002.

[8] M. Harré, T. Bossomaier, A. Gillett, and A. Snyder, "The aggregate complexity of decisions in the game of go," *European Physical Journal B **ERA A***, vol. in press, 2011.

[9] R. Wicks, S. Chapman, and R. Dendy, "Mutual information as a tool for identifying phase transitions in dynamical complex systems with limited data," *Phys. Rev. E*, vol. 75, 2007.

[10] S.-J. Gu, C.-P. Sun, and H.-Q. Lin, "Universal role of correlation entropy in critical phenomena," *journal of physics A*, 5 2006.

[11] F. Gobet and P. Chassy, "Expertise and intuition: a tale of three theories," *Minds and Machines*, vol. 19, pp. 151–180, 2009.

[12] W. Chase and H. Simon, "The mind's eye in chess," in *Visual Information Processing*, C. W.G., Ed. Academic Press, NY, 1973, pp. 215–281.

[13] F. Gobet and H. Simon, "Five seconds or sixty? presentation time in expert memory," *Cognitive Science*, vol. 24, pp. 651–682, 2000.

[14] J. Rubin and I. Watson, "A memory-based approach to two-player texas hold'em," in *AI 2009: Advances in Artificial Intelligence, Proceedings*, ser. Lecture Notes in Artificial Intelligence, A. Nicholson and X. Li, Eds. Springer, 2009, vol. 5866, pp. 465–474, 22nd Australian Joint Conference on Artificial Intelligence DEC 01-04, 2009 Melbourne, AUSTRALIA.

[15] J. Tromp and G. Farnebäck, "Combinatorics of go," *Computers and Games*, pp. 84–99, 2007.

[16] S. Gelly and Y. Wang, "Exploration exploitation in go: Uct for monte-carlo go," in *Twentieth Annual Conference on Neural Information Processing Systems (NIPS 2006)*. Citeseer, 2006.

[17] D. Stern, R. Herbrich, and T. Graepel, "Bayesian pattern ranking for move prediction in the game of go," in *Proceedings of the 23rd international conference on Machine learning*. ACM, 2006, pp. 873–880.

[18] C. Shannon and W. Weaver, *The Mathematical Theory of Communication*. Univ. Ill. Press, Urbana, 1949.