Seeing the Big Picture: Influence of Global Factors on Local Decisions

Terry Bossomaier*, Michael Harré*†, Vaenthan Thiruvarudchelvan*
*Centre for Research in Complex Systems
Charles Sturt University
Bathurst, Australia
e-mail: {tbossomaier,vthiru}@csu.edu.au
†Centre for the Mind
University of Sydney
Sydney, Australia
e-mail: michael.harre@sydney.edu.au

***Abstract*—This paper extends recent work studying the development of human expertise in the game of Go. Although it appears like a simple game on the surface, Go is actually the most difficult of all established games for artificial intelligence, with no computer program yet reaching the top international level on a full 19×19 board. On smaller boards with sizes like 9×9, computers *are* competitive, implying that the understanding of complex global interactions is the key to human superiority. The temporal analysis of game positions yields some interesting insights in into local/global analysis. By mining thousands of positions from online games, 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. Using perceptual templates introduced in other work, we demonstrate that this global context appears at the very earliest stages of cognition.**

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

I. Introduction

The big picture often influences or overrides 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, using data mined from thousands of decisions in online games. 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. This contrasts with Chess, where the IBM computer Deep blue [6][7] was able to defeat world champion Garry Kasparov.

We also demonstrated therein, from calculation of mutual information (eqn. 5) between moves, that one of these has the character of a *phase transition* [8]. The idea of a phase transition comes originally from physics, from the study of phenomena like 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 adding edges to random graphs. At a certain point each graph shows a transition: the average path length (the number of steps from one node to another) rises to a peak, and then drops back down again.

A dynamical system example 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 partially align. Phase transitions occur in this system as particles flow around in groups, like flocks of birds, but dynamically—continually forming and dissolving.

Mutual information is a precisely defined quantity, originating from Shannon's mathematical theory of communication [10]. It 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 above, 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][11] and, along with other characteristics like long-range order and power-law behaviour, is thought to be a general property of phase transitions.

Previous work [8] has already demonstrated phase transitions in collective human decisions in Go. In this paper, 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 Professionals [8]. We also present evidence that there is global influence on local decisions, and that the influence is greatest during the phase transition. The evidence for the global factors arises from temporal analysis: the next move is more predictable given the sequence which led up to it, compared with just using the position at which it is made. We argue below that this arises from the global information inferred from the sequence.

Section II describes the conceptual background or expertise (sections II-B and III-B discuss *perceptual templates,* which form one of the earliest stages of processing of a Go board. It turns out that for professional players, these templates have a strong non-local character, supporting the findings from mutual information.) Sections III and IV describe the methods and results respectively. The discussion (section V) and conclusions (section VI) round off the paper.

II. Expertise and Perception

The study of expertise in games owes much to Fernand Gobet and his colleagues, summarised in his book *Moves in Mind* [1]. However, the methods used in Go in this paper rely on a new methodology introduced in [12]. The next two sections discuss these in turn.

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][13]. One of the key ideas, essentially from Nobel Laureate Herbert Simon, is that human expertise involves building a huge library of patterns [14][15]. The application of these ideas in artificial intelligence for games is relatively new however [16].

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

Further recent advances have been limited, particularly in Go, where a combination of the gamespace complexity [17] and a lack of genuinely human-like heuristics like 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 [18], some progress has been made in achieving strong amateur play. However, these techniques do not address the inherent complexity of the game nor the techniques that humans have developed in order to address this, almost completely because it is difficult to investigate.

Of relevance to game players is the current state of the game, the likely future states of the game and in what order those future moves will be played. The current state of the game is very well approximated by the pieces currently on the board (this excludes some technical rules about repeated positions that are only rarely relevant), and these can be divided up loosely into tactical, strategic and distracting pieces. Tactical pieces are involved in local battles for territory, while strategic pieces play a role in long-term plans spanning the entire board. Distractors play only weak roles in either of the preceding plan types. Of course, a single stone can participate in both local and global strategies. In terms of future states of the game, we considered only local patterns and what was played in the local area – a purely tactical aspect of the game. This leaves only the strategic relationships as a source of information that might perturb the actual moves made. It is this external influence on tactical plays that is implicit in the global contextual analysis of this paper.

We argue that the sources of information players use in order to make good decisions are of two types: *local* and *global.* Every level of player in our study has learned a great deal about the game of Go over the course of their lives; we now want to make explicit and quantify this information. We do this by looking at the probability distributions of moves made in a variety of different positions. The relevance of the division of the problem space into these two parts can be seen in the work of Stern et al. [19]. 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 a Bayesian framework for matching moves to positions.

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 had not been reported in the research literature before our recent work.

B. Kohonen Maps & Perceptual Templates

If local decisions involve global factors, the question arises as to where in the cognitive hierarchy global information appears. We use the recent work on *perceptual templates* to show that it starts at the very lowest levels. Perceptual templates are the building blocks of perception, experienced preattentively and fundamental to rapid decision making – the instant appraisal of situations by experts, the guiding of eye movements and expert memory for real-world positions.

A novel way to determine such perceptual templates involves the use of Kohonen maps trained on game data [12]. The templates so found can then be analysed for global properties. Teuvo Kohonen [20] introduced self-organising maps (SOMs) as a model of human visual information processing. Although they help explain some structural characteristics of the visual cortex, they have found considerable practical use in the signal-processing domain, especially image processing.

A SOM is a competitive learning process, comprising of a selectable number of neurons. Each neuron has a random weight vector, and a set of inputs of the same dimension. In the case of an image, the inputs would be the colour components of each pixel.

Training proceeds as follows. A pattern from the training set is presented to each neuron in the map. There will be one neuron which is closer (different metrics of proximity may be used) to the pattern than any of the others. The weights of

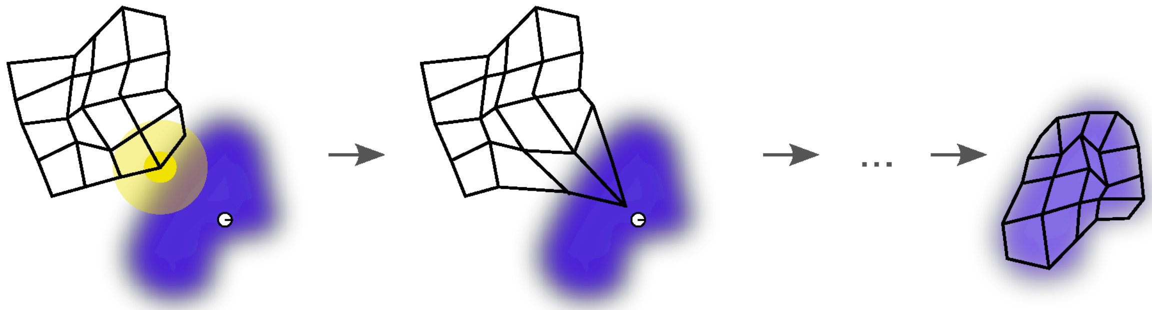

Figure 1. This illustration captures the essence of the Kohonen map or Self-Organising Map (SOM), mapping two dimensions to two. The blue cloud represents the training dataset, and the intersections of the grid represent the weight vectors of the neurons in the map. Initially the weights are random and the map does not correlate with the data (left). Training the map involves finding the neuron with the closest weight vector (yellow dot) for each datapoint (white dot), and adjusting the weight towards it (center). After repeated training, the map comes to approximate the dataset (right). In our work, the dataset and weight vectors are 361-dimensional (intersections on the Go board), while the map is 2D with 50x50 neurons. Copyright Mcld [CC-BY-SA-3.0 or GFDL], via Wikimedia Commons. *http://commons.wikimedia.org/wiki/File:Somtraining.svg*

this and neighboring neurons are then increased while neurons further away are decreased. Presentation of patterns is repeated until the weights converge. Fig. 1 illustrates this process.

The competitive, winner-takes-all strategy is difficult to analyse theoretically. There are some results on the proportion of neurons devoted to the probability density of particular features, but such results are often derived under simplifying assumptions. Empirically what happens is that the neurons end up representing commonly occurring sub-features. Unlike multi-layer feedforward networks, the basic SOM has no capacity for translational or geometric transformation. In image processing this is often a disadvantage, but with the fixed dimensions of the Go board, it is actually advantageous here.

We have used two different techniques in this work that have quite separate roles and interpretations in making decisions in any complex environment. The perceptual templates are a learned, static and sparse representation of the game state that is implemented as a perceptual guide to inform higher order cognitive processes. One type of high-order cognitive processing is exact pattern matching and planning, and it is this aspect that is studied using information theory. These two aspects are necessarily disjoint as they are plausible subsystems that are likely present in Kahneman's 'Dual Processing' account of human cognition [21]. In this interpretation, perceptual processes are fast to implement (as they are parallel processes), slow to learn and unlimited in their capacity. In contrast, planning and evaluation processes are most likely slow to implement (as they are serial processes) but are deployed dynamically and have limited capacity.

III. Methods

Harré and Bossomaier [8] examined game trees six moves deep (i.e., three black and three white) 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 Professional (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 7×7 board sections in the corners, 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 six moves played in the region, they are not necessarily the first six moves of the game.

A. Information Theory

Information theory was introduced by Shannon [10] to study the communication of data through channels. It has since proved immensely useful across many domains. It is essentially a statistical concept, quantifying the gains and losses from choices amongst alternatives.

Shannon defined the information of an event in bits, η_x as:

$$\eta_x = -\log_2 P(x) \quad (1)$$

where $P(x)$ is the probability of observing the event x. If we now average over the set M, of all events which could occur, denoted by x_i, we get the average information, which is also known as the *entropy*:

$$H(M) = -\sum_i P(x_i) \log_2 [P(x_i)] \quad (2)$$

Applying this to moves in Go gives us a measure of how likely one move is over another. If the entropy is zero, then $P(x_i)$ must be one for some x_i and zero for all others. Thus the move is completely determined without ambiguity. Entropy is a measure of disorder or randomness and is maximal when the probabilities for all events are the same.

We need one additional concept, the *conditional entropy*. The likelihood of any given move is of course dependent upon previous moves up to that point. Thus we can ask what the probability of a move is in a given position. From this, we can ask what the entropy is in the set of moves if we know the position in which the move was made. If we now average over all positions, we get the entropy of moves given positions. This is the conditional entropy, defined in eqn. 4.

For each possible move on the Go board m_i, three probability distributions were computed:

1) the probability of the move occurring, $P(m_i)$
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 (eqn. 5) were calculated, but this paper addresses findings from the entropies alone. A discussion of the primary results from mutual information is given in [8].

The move entropy, $H(M)$, is taken over all moves which can arise at each level in the game tree (i.e., for the six moves in the sequence):

$$H(M) = -\sum_i P(m_i) \log_2[P(m_i)] \tag{3}$$

For the first move in the region there are 49 possible positions, decreasing to 44 after five moves, giving a maximal entropy of $\log_2 44 = 5.5$ bits, which would occur if all moves were equally likely. But since the moves are chosen strategically, they are far from random, so the measured entropies are much lower than this.

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

$$C(M|Q) = -\sum_i \sum_j P(m_i|q_j) \log_2[P(m_i|q_j)] \tag{4}$$

From the conditional entropies, we can calculate the mutual information, $I(Q, M)$ using Shannon's formula [10], eqn. 5:

$$I(Q, M) = H(M) - C(M|Q) \tag{5}$$

The same expressions are used for an ordered sequence of moves, replacing q_j with s_j. These entropic quantities are now calculated across all ranks, from 2 Kyu Amateur (*am2kyu*), through the amateur ranks to *am6d*, onto the highest rank of 9 Dan Professional (*pr9d*). The results are shown in Figs. 1–3.

B. Perceptual Templates

For our work, we use a separate 50x50 SOM (2500 neurons) for each of the 361 intersections of the Go board. Each SOM is trained on board states directly preceding a move at that point, mined from the online gameplay database. Board states are represented as linearized length-361 vectors with values equal to -1, 0 or 1, representing black stone, empty or white stone respectively. Games were normalized to always start with a white stone, and no deduplication along axes of symmetry was carried out. The weight vectors of the neurons was also constrained to this range, facilitating easy template extraction.

Threshold	**# Templates**	**Average Size**
0.9	10,929	11.1
0.8	26,318	13.8
0.7	55,553	15.2
0.6	145,534	16.5
0.5	364,557	18.3

TABLE I. Number of perceptual templates and their average sizes (maximum Euclidean distances) per threshold.

Further details may be found in [12], from which trained SOMs were reused. The ones used for this paper are taken from games 5 Dan Professional and above. 18,000 games were used in training each map.

The spatial topology of a trained SOM is usually of significance in typical uses, however we discard this information. We consider the weight vectors of each neuron at every point as potential perceptual templates. Neurons which have strongly learned patterns of stones across the board will then be extracted as templates. Thus the number of potential templates equals the total number of neurons, $361 \times 2500 = 902,500$.

However, since there are on average only about 7 training games per neuron, the learned weights are still quite noisy. Therefore, actual templates are extracted from the weight vectors by thresholding the weights to 1, 0, or -1 using thresholds of k and $-k$, for values of k of 0.5, 0.6, 0.7, 0.8 and 0.9. After thresholding, empty and duplicate templates are removed, leaving useful templates. Table I records the resulting number of templates at each threshold. Examples of these templates are shown in Fig. 2, which may be locally clustered as in subfigures a. and b., but are often non-local as in c. and d. Further procedural details can be found in [12].

IV. Results

Fig. 3 summarises the key findings of the paper. It shows the conditional entropy as a function of move in the sequence of six, averaged across all ranks, both amateur and professional. Note that the moves are logged as they occur in the game. They are *not* necessarily in sequence. In other words, this is not a game on a small board region but a window on a full 19×19 game. Since the standard of play is professional for this analysis, extremely weak moves are unlikely to occur and will not appear in the game records. Error bars are calculated as in Harré et al. [5]. Up to move three, 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.

Fig. 4 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 (Section III). The entropies observed are of course much lower – usually less than 2 bits – reflecting the structure inherent in the game.

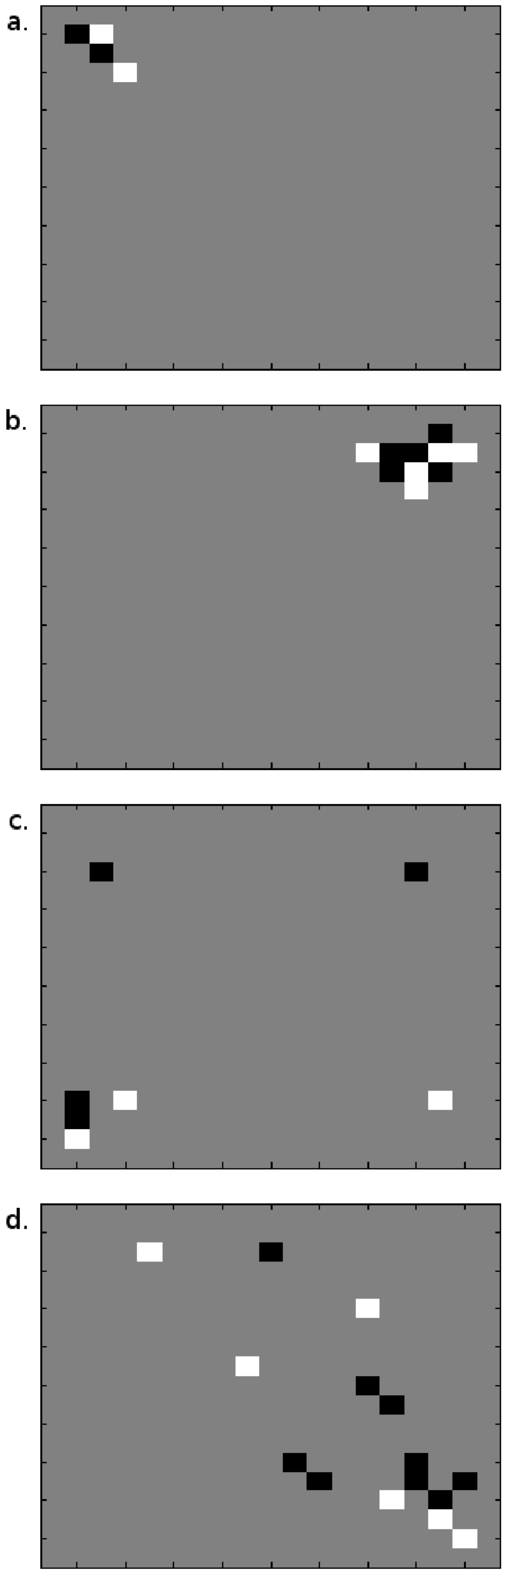

Figure 2. Four templates arising from a threshold of 0.9. Most templates are clustered in a corner of the board like a. and b. Template c. spans multiple corners, as games commonly develop. Template d. reflects a game which spreads outwards rather than clustering in the corners.

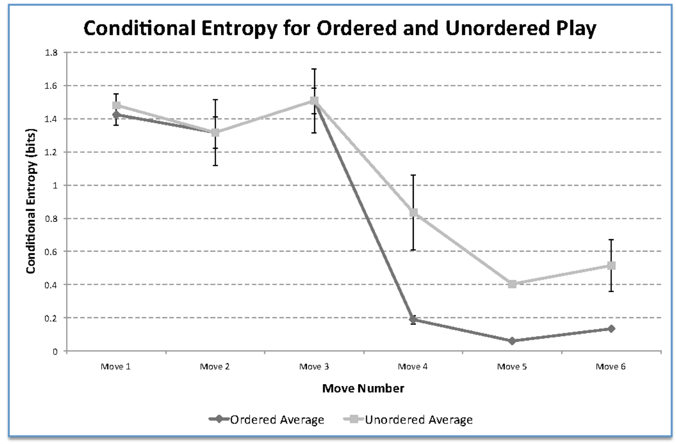

Figure 3. Conditional entropy (eqn. 4) as a function of move averaged over all ranks

The entropy for the third move is on the order of 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 of stones on the board increases. The entropy summed over all moves declines linearly to the maximum amateur rank and then increases gradually from from the first professional rank onwards.

Fig. 5 shows the entropies resulting from positions which arose from a particular sequence of play. These entropies are around 3 bits smaller than 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 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.

Turning to the templates derived from the SOMs, as the threshold is reduced from 0.9 to 0.5 (Figs. 6–10) the number of templates increases exponentially, but remarkably, the shapes of the histograms do not change so rapidly. So the number of stones in a template is around 5 for the bulk of them, for thresholds 0.6 and upwards. As the threshold drops, the new templates which appear are new patterns of a small number of stones, rather than patterns with increasing numbers of stones. (The last bar includes all templates with higher stone counts).

Table I includes the averages of the maximum Euclidean distance between two stones in each template, per threshold. We see that it rises fairly linearly with threshold. Figs. 11 – 15 are histograms of those distances.

The theoretical maximum distance is 27, but since this includes stones on opposing corners, the maximum observed distance will usually be lower. At a threshold of 0.9, over half

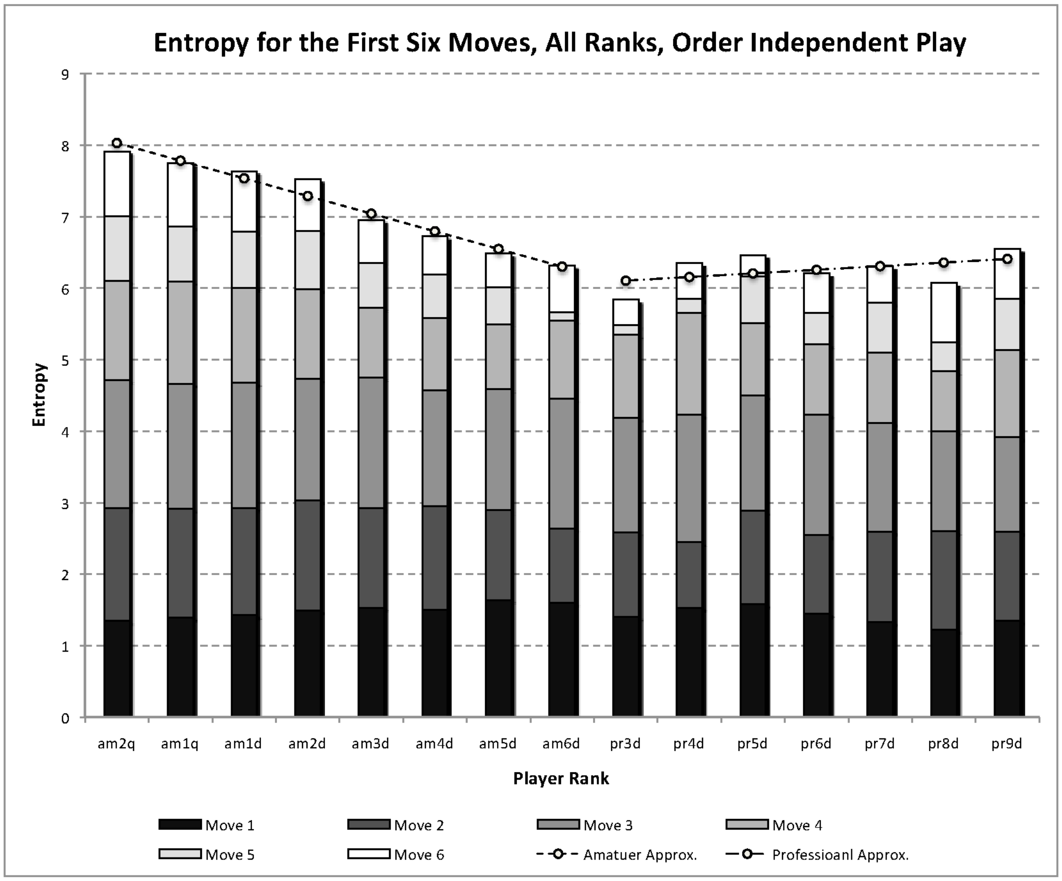

Figure 4. Entropies for moves from a given board position (reprinted from [8])

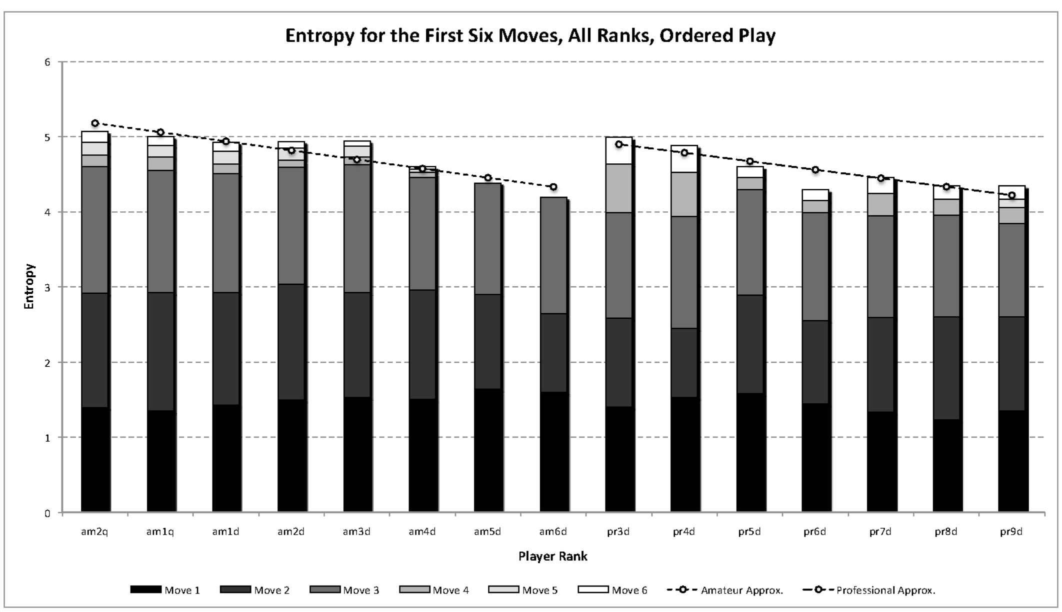

Figure 5. Entropy for the first six moves shown as a stacked bar chart. The black bars represent the entropy at move 1, the dark grey at move 2 and so on for all six moves. The dashed regression lines show the total entropy for the amateur and professional sequences.

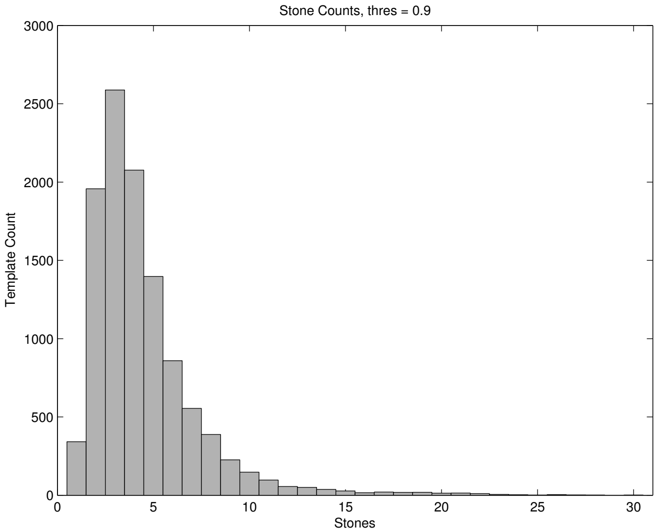

Figure 6. Number of templates containing a given number of stones, threshold = 0.9.

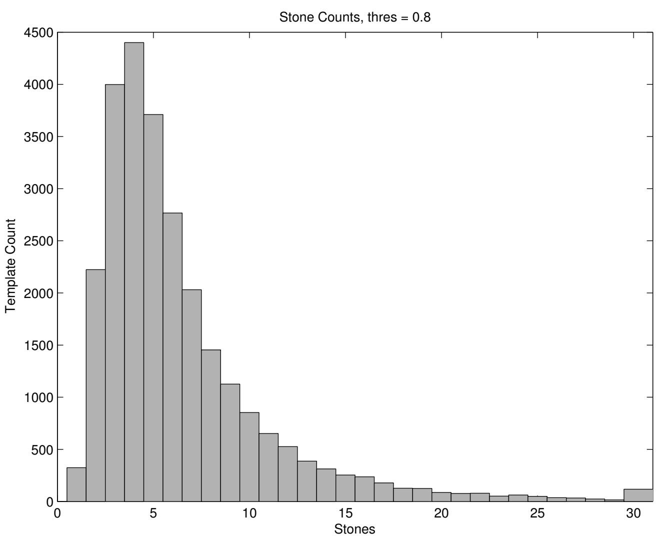

Figure 7. Number of templates containing a given number of stones, threshold = 0.8.

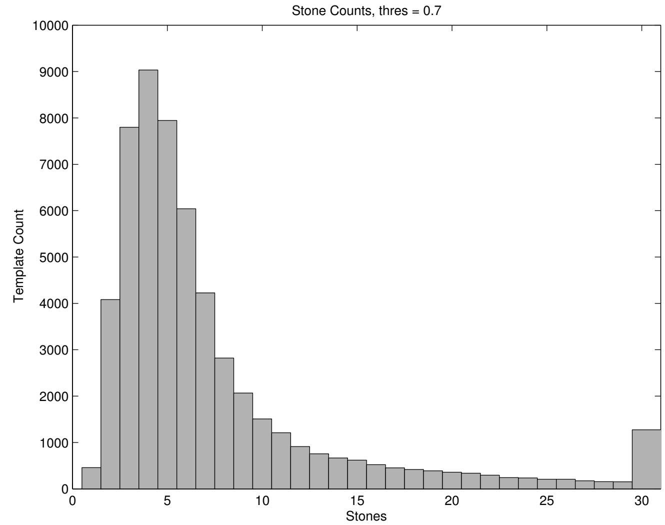

Figure 8. Number of templates containing a given number of stones, threshold = 0.7.

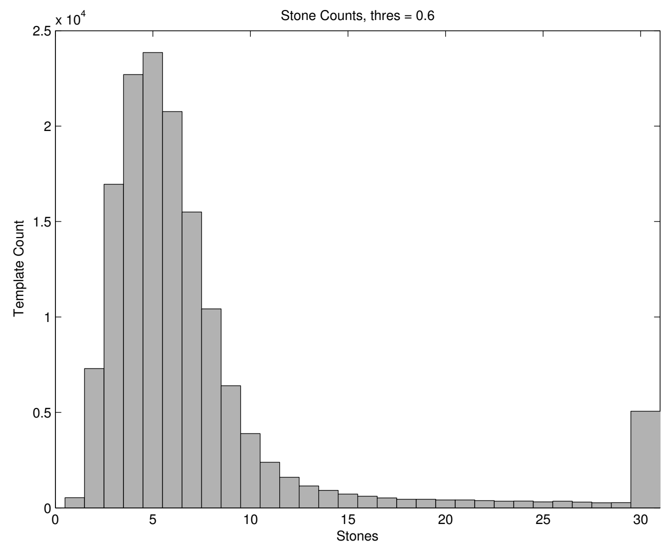

Figure 9. Number of templates containing a given number of stones, threshold = 0.6.

the templates have a distance of 12 or greater, implying that they include stones in different corners of the board. There could be a small number of templates occupying one corner and the centre, but these would need to have distances around 12, where we find a minima instead.

Also salient are the twin peaks around 3 and 18, possibly representing corner clusters and cross-board patterns respectively. As we lower the threshold, the proportion of templates spanning the board increases until it dominates at 0.5, while the peak at 3 stones vanishes.

V. Discussion

There are three very interesting features of these results, which we consider in turn: a) the difference between ordered and unordered play, b) the way the conditional entropy varies with rank, and c) how perceptual templates span the entire board.

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 relationships between localised groups of stones 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 next three moves.

The gradual decline in entropy with rank for amateur and professional reflects a gradual reduction in the space or range

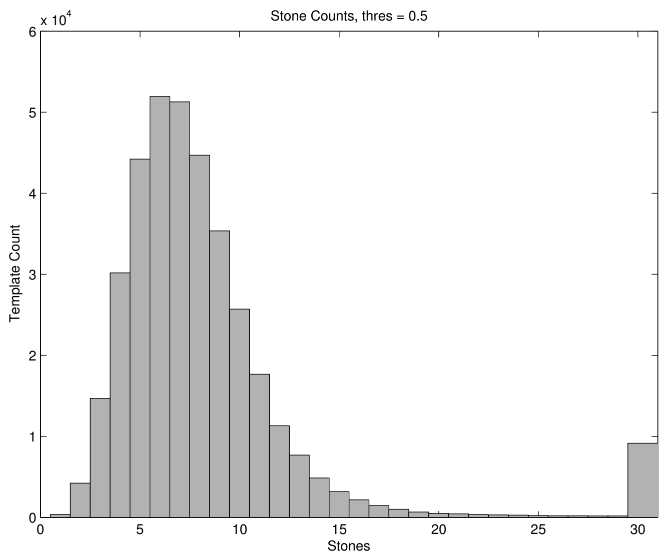

Figure 10. Number of templates containing a given number of stones, threshold = 0.5.

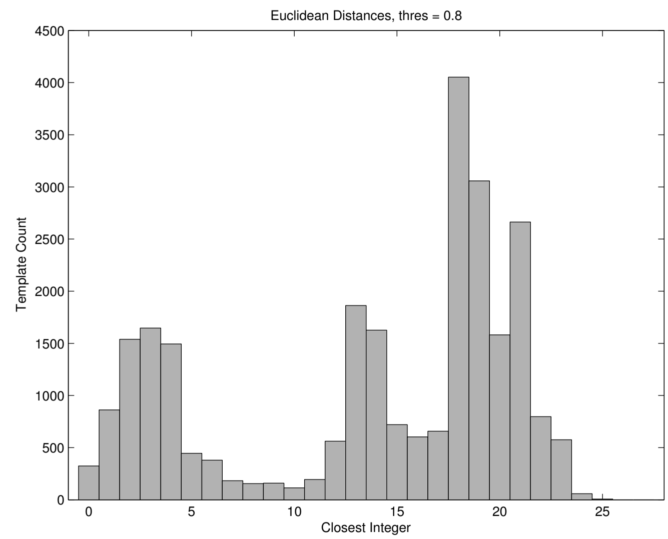

Figure 12. Maximum Euclidean distance between stones in a template, threshold = 0.8.

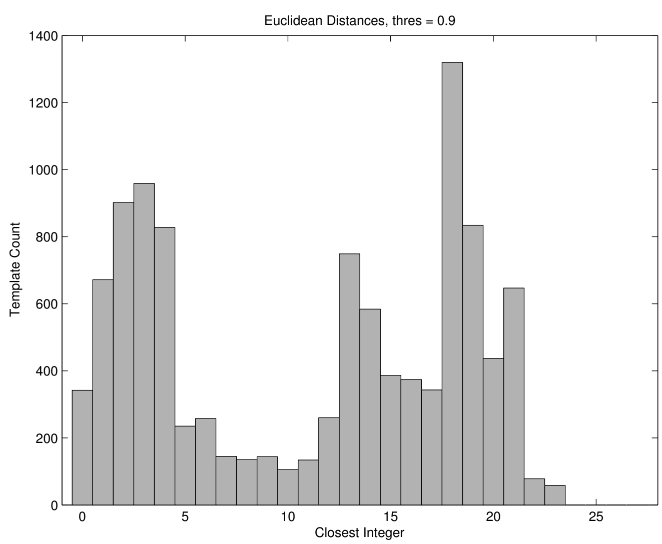

Figure 11. Maximum Euclidean distance between stones in a template, threshold = 0.9.

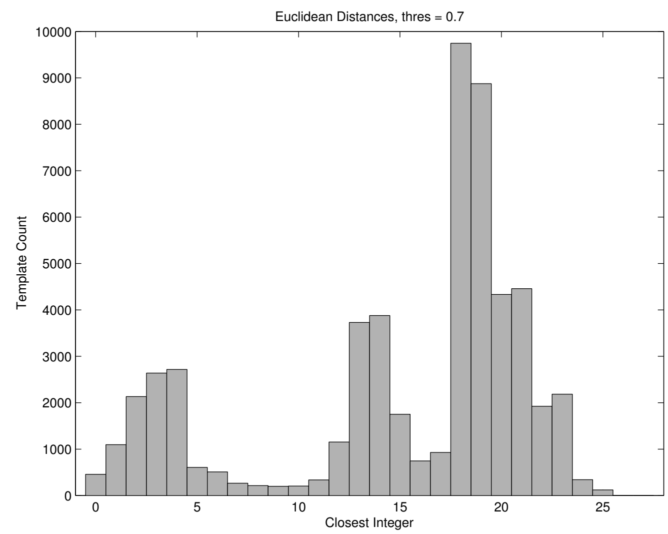

Figure 13. Maximum Euclidean distance between stones in a template, threshold = 0.7.

of options, which we could conceive of as the elimination of poor moves in established situations, similar to mastering the openings 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 [10].

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 opponent-reading skills in a complete information game such as Go, we believe that it is relatively insignificant. The strategical influence of the other stones on the board that were not within the local area of study, is a much more significant factor. The changing influence that non-local information has on decisions during a game, is evident in the significant drop-off in the conditional entropy after move three in Fig. 3, a drop in shared information of nearly an order of magnitude for the ordered play and about half that for unordered play. This is consistent with the observation that, at the time of writing, the best computer Go programs are close to professional-level on small boards like 7×7, but rapidly deteriorate on larger boards, as global influences become important.

This change in conditional entropy in the corner regions of

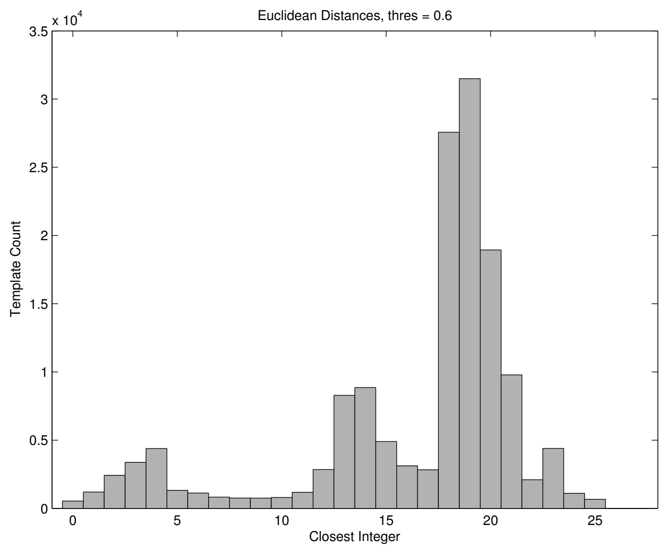

Figure 14. Maximum Euclidean distance between stones in a template, threshold = 0.6.

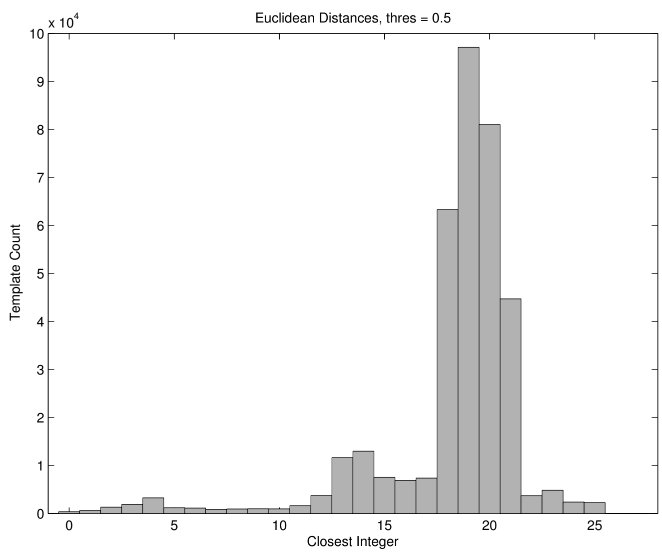

Figure 15. Maximum Euclidean distance between stones in a template, threshold = 0.5.

the board as the game progresses might be due to the shrinking size of the move space as the board fills up. While 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 1/43 per move. It is also possible, but exceptionally unlikely, that after the third move, players choose much more randomly – i.e., without concern for pieces on the board, local or not – than they did for the first three moves. Considering the vast training literature available to players that readily teaches them the many different variations of the first six moves within a corner, and how to contextualise these decisions by considering what stones occupy nearby areas, we consider this to be an unlikely proposition.

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 conditional entropy. This is perhaps not so surprising when considered in the light of the state of the game itself, after three moves have been played in the corner. These first moves can be thought of as establishing the board layout in terms of an 'opening book'; highly stylised placement of local stones, where the local pattern can be thought of as effectively uncoupled from the rest of the board, or at least equally coupled for these first moves. This coupling then changes significantly from the fourth move onwards, where greater consideration needs to be afforded to the other pieces on the board. This change in the focus of gameplay significantly reduces the information coupling between local moves and local stones on the board.

The use of global information is supported by the analysis of the perceptual templates. A large fraction of templates cover more than one corner of the board, implying that global analysis *starts at the very earliest perceptual levels*. As the threshold is reduced to 0.5, the number of local templates actually gets eclipsed by non-local ones. An additional finding from our results is that the majority of perceptual templates contain less than ten stones, regardless of noise threshold and number of templates. Even at the 0.5 threshold, where most templates are non-local, the average number of stones remains small. This median figure of around 5 – 7 is on the order of human working memory capacity and similar to the figures in Gobet's CHREST models [1]. This is therefore an important issue, subject to future research.

The complete disappearance of entropy at the high amateur ranks is very interesting. It suggests that 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 by Harré et al. [8]. They observed a peak in mutual information at the transition to professional play, 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 the phase transitions might be predicted.

A. Implications for Computer Go

The objective of this study was to determine some characteristics of human Go expertise. These may subsequently be fed into the computer Go domain, but that was not our motivation here. Our analysis is once-off, so the time-complexity of our computations is irrelevant. The methods as described in this paper have never been used before, and no prior work has attempted to identify the influence of global factors in Go.

VI. Conclusions

The analysis of large volumes of data has generated powerful new insights into human cognition in the Game of Go, with potential applicability to other domains. We have shown that low-level perceptual templates of professional players are non-local, i.e., include features from the whole board. The paper links this to earlier work on mutual information in local positions, which we infer to be influenced by global factors. The *sequence* of moves leading to a position was shown to provide more information about the next move than the position alone, which could be accounted for by global contextual information provided by the former.

The big challenge for future work is to determine if these properties hold in other domains. Poker is the ideal next game to study: it is the second most difficult established game for computers to play well, and has the additional features of incomplete information, stochastic elements and theory of mind.

In 2012, Zen, one of the top computer Go programs, won games against 6 Dan Amateur players, and came much closer to beating professional-level players than any program has before. But computer Go relies heavily on Monte Carlo Tree Search, which is nothing like human tactics or strategy. It remains desirable to try to understand and mimic the way humans learn and play. A big open question is whether the future of game playing software, or software in general, will adopt these strategies. The human brain trades off search speed and accuracy for robustness and possibly scalability. Human decisions may sometimes be inferior, but they rarely exhibit the catastrophic failures resulting from software bugs. The extent to which the strategies of human expertise and computer algorithms hybdridise will be one of the really exciting topics of the next decade.

ACKNOWLEDGEMENT

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

REFERENCES

[1] F. Gobet, *Moves in Mind; The Psychology of Board Games*. Psychology Press, Sep. 2004. [Online]. Available: http://www.worldcat.org/isbn/1841693367

[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, and A. Snyder, "The development of human expertise in a complex environment," *Minds Mach.*, vol. 21, pp. 449–464, August 2011. [Online]. Available: http://dx.doi.org/10.1007/s11023-011-9247-x

[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] Harré, M. S., Bossomaier, T., Gillett, A., and Snyder, A., "The aggregate complexity of decisions in the game of go," *Eur. Phys. J. B*, vol. 80, no. 4, pp. 555–563, 2011. [Online]. Available: http://dx.doi.org/10.1140/epjb/e2011-10905-8

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

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

[12] M. Harré and A. Snyder, "Intuitive expertise and perceptual templates," *Minds and Machines*, pp. 1–16, 2011, 10.1007/s11023-011-9264-9. [Online]. Available: http://dx.doi.org/10.1007/s11023-011-9264-9

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

[14] 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.

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

[16] 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.

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

[18] 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.

[19] 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.

[20] T. Kohonen, "Self-organized formation of topologically correct feature maps," *Biological Cybernetics*, vol. 43, pp. 59–69, 1982, 10.1007/BF00337288. [Online]. Available: http://dx.doi.org/10.1007/BF00337288

[21] D. Kahneman, *Thinking, Fast and Slow*. Farrar, Straus and Giroux, 2011.