Strategic Information in the Game of Go

Michael Harré , Terry Bossomaier, Ranqing Chu, Allan Snyder

Abstract—We introduce a novel approach to measuring how humans learn based on techniques from information theory and apply it to the oriental game of Go. We show that the total amount of information observable in human strategies, called the strategic information, remains constant for populations of players of differing skill levels for well studied patterns of play. This is despite the very large amount of knowledge required to progress from the recreational players at one end of our spectrum to the very best and most experienced players in the world at the other and is in contrast to the idea that having more knowledge might imply more 'certainty' in what move to play next. We show this is true for very local up to medium sized board patterns, across a variety of different moves using 80,000 game records. Consequences for theoretical and practical AI are outlined.

Keywords—Board Games, Cognitive Capacity, Decision Theory, Information Theory.

I. INTRODUCTION

The game of Go has been played in the Asian region for more than 2,500 years and has a prominent place within the culture of this region. In recent years it has become one of the most popular board games in the Orient where there are world titles, professional players as well as thriving online communities. There are also competitive matches between the various artificial intelligence (AI) programs in order to decide which AI is the best in the world. In this respect Go has taken up the mantel from Chess as one of the grand challenges of AI [1] as well as providing a rich source of information for the study of human cognitive behaviour in psychology [2] and neuroscience [3], [4]. From these recent studies and from reviews such as Burmeister and Wiles' work [5] Go will form a cornerstone problem for practical applications of AI and expert systems.

II. THE GAME OF GO AND AI

A. The Game of Go

Go is usually played on a square board composed of a 19×19 grid pattern although 7×7, 9×9 and other sizes are less frequently used. Each player has a designated set of coloured 'stones', either black or white, and play begins with black placing a stone on one of the intersections of the board. Play is taken in turns until both players pass consecutively at which point the game stops and the board is scored. The objective is to capture more territory than your opponent by surrounding territory with your stones, see[1] for a basic introduction to game and its rules.

Only a small number of additional concepts are required for our work. First, all players that will be of interest are ranked. Amateur ranks who are beginner to an intermediate level of play range from 30 kyu - 1 kyu, 30 kyu being the lowest and 1 kyu being the highest in rank. There are a further eight amateur ranks at the 'Dan' level, 1-Dan through to 8-Dan in ascending order of skill. There are also the ranks 1-Dan through 9-Dan of the professional players, the next step after the amateur ranks. A player becomes a professional by receiving a diploma through one of the professional Go associations. The competition is extremely high for this certification and there are comparatively very few westerners who have achieved a professional ranking.

B. Recent Developments in Go AI

Up until very recently it was possible for a human player with much less than a years experience to easily beat the worlds best computer software at the game. However, recent advances in brute force search algorithms has lead to significant improvements in this area.

In particular, the software program MoGo has recently won against several professional players. In 2008 MoGo won one game out of three against Catalin Taranu, a 5-Dan professional on a 9×9 board. MoGo lost a 19×19 game against the same opponent despite a nine stone handicap. Later that year MoGo won against Myungwan Kim, an 8-Dan professional, on a 19×19 board with a nine stone handicap. This second competition was run using a much more powerful computer than the first, the Huygens supercomputer provided by SARA and NCF[2]. MoGo has also won against Zhou Junxun, a 9-Dan professional, in a seven stone handicap game[3]. Other successful computer Go players include "The Many Faces of Go" and "Crazy Stone", see[4] and references therein for the most recent results.

Results such as these have inspired recent comments by Feng-Hsiung Hsu, the principal designer of Deep Blue which in 1997 defeated then world champion Garry Kasparov at chess. He as suggested that Go will fall to 'hard AI', just as Chess did, within the next ten years[5].

Supposing this is an accurate estimate, there is then a race on to see which of two possible strategies can first address AI expertise in the game of Go. One is the direct tree search method and the recent successful variations thereof [6], [7] which involve little or no modelling of human strategies. The alternative is the development of psychologically based models which emphasises both the strengths and weaknesses

of humans [8]. It is this second issue to which this work is targeted, emphasising the observable limitations of human play and how that might inform AI models of expert players.

Of particular interest to the AI and cognitive sciences communities is the complexity of Go. This is especially important due to the ability of humans to be able to reduce this complexity to a manageable level while still being capable of defeating the very best AIs at the top level. Chess [9] has a state space complexity of 10^{43} and a game tree complexity of 10^{120}. In comparison, Go [10] has a state space complexity of 10^{171} and a game tree complexity of 10^{360}. How humans are able to reduce this complexity, and how this evolves as a player's rank increases, is a fundamental question for the development of artificial learning systems that are able to develop and eventually behave more like humans.

III. THEORETICAL FRAMEWORK

Our method of analysis is statistical in nature. We are interested in the observed probability distributions over move choices for given board configurations. The stochasticity has two distinct contributions, the psychological uncertainty of each individual player and the stochasticity that arises from sampling from a population. Accepting the latter as an understood property of population sampling, we briefly discuss the former.

Psychologically players are influenced by mood, temperament and biological imperatives such as thirst and hunger. At a more pragmatic level, position assessment, tactical play, inferences about the other player's strategies all play their part in influencing how individuals choose their next moves. It is also true in principle that for every board configuration there is an optimal move. However people do not exhaustively search all possible alternatives for the perfect move. So despite the fact that Go is a game of perfect information, in practice there is always some uncertainty as to whether any particular move is the 'perfect' move. These considerations lead us to believe that there will always be a naturally stochastic nature to decisions made in games such as Go.

The methods developed in this work do not distinguish between these two sources of noise in player's choice of strategy. However we are looking to understand how a population of players behave, not necessarily the individuals. With this in mind we next introduce how we measure the difference in strategy choice between different populations of players.

A. An Entropic Measure of Skill

First we introduce some notation. We divide the Go board up into $i \in \{1, \ldots, n\}$ regions where the region defined as the whole board has the unique label r. The union of all r_i is a complete cover of the board, i.e. $\bigcup_i r_i = r$, however it is not necessarily true that the intersection of all regions is an empty set. Given the i^{th} region of the board r_i, the set of legal moves available in that region belong to a set Ω_i where $\sigma_j \in \Omega_i$ is one of the possible legal choices of moves in r_i. We reserve σ to denote a discrete stochastic variable and σ_j a specific outcome of a random event, i.e. the probability mass function is $p(\sigma_j) = \Pr\{\sigma = \sigma_j\}$, $\sigma_j \in \Omega_i$, $\sum_j p(\sigma_j) = 1$, $p(\sigma_j) \geq$

$0 \ \forall \ j$. By a move made in a given r_i we mean the next stone to appear in r_i, regardless of how many moves have since been played elsewhere on the board outside of r_i.

We are concerned with different distributions over the same Ω_i called $p(\sigma)$. Given an arbitrary function $g : \Omega_i \rightarrow I\!R$ the expectation of $g(\sigma)$ for a known $p(\sigma)$ is: $E_p[g(\sigma)] = \sum_{\sigma_j \in \Omega_i} p(\sigma_j)g(\sigma_j)$. For the specific case of $g(\sigma) = -\ln(p(\sigma))$ we define the entropy as [11]:

$$
\begin{aligned}
H(p(\sigma)) &= E_p[-\ln(p(\sigma))], && (1) \\
&= -\sum_{\sigma_j \in \Omega_i} p(\sigma_j)\ln(p(\sigma_j)) && (2)
\end{aligned}
$$

Grünwald and Dawid [12] call $-\ln(p(\sigma_i))$ the log-loss function. Furthermore, given a second probability distribution $m(\sigma)$ and we take the expectation of $g(\sigma) = \ln\frac{p(\sigma)}{m(\sigma)}$ we get the relative entropy between $p(\sigma)$ and $m(\sigma)$ defined as:

$$
D(p(\sigma)\,\|\,m(\sigma)) = E_p\left[\ln\frac{p(\sigma)}{m(\sigma)}\right] \qquad (3)
$$

The relative entropy is a measure of the divergence of $p(\sigma)$ from $m(\sigma)$ and it is this divergence which we shall use to measure the differences between players of different ranks. It has been known since its inception [13] that this divergence measures the (average) amount of information required to discriminate $p(\sigma)$ from $m(\sigma)$.

There are properties of the relative entropy that should be remarked upon. If $p(\sigma) = m(\sigma)$ the relative entropy is zero and therefore the distributions are indistinguishable. For any $p(\sigma) \neq m(\sigma)$ we have: $E_p\left[\ln\frac{p(\sigma)}{m(\sigma)}\right] > 0$, that is for fixed $m(\sigma)$, $D(p(\sigma)\,\|\,m(\sigma))$ is minimised when $p(\sigma) = m(\sigma)$. In general $m(\sigma)$ will be taken to be our reference distribution, it is the distribution over moves played by the best players in our data-set, and so $D(p(\sigma)\,\|\,m(\sigma))$ measures the divergence of distribution $p(\sigma)$ from that of the best players, $m(\sigma)$.

In our data we are confronted with having potentially infinite values for $D(p(\sigma)\|m(\sigma))$ when $m(\sigma_i) = 0$. To address this we adjust equation 3 in the following way:

$$
D_s(p(\sigma)\|m(\sigma)) = E_p\left[\ln\frac{p(\sigma)}{0.5(m(\sigma) + p(\sigma))}\right] \quad (4)
$$

This is called the Jensen-Shannon Divergence [14]. For the purposes of this work, this alternative formulation maintains the useful properties of equation 3 while ensuring the divergence is always finite.

The principle of measuring relative entropy and its relationship to player rank can be made more precise. Suppose we are given a reference distribution $p_\infty(\sigma)$ representing an all knowing oracle of 'infinite' skill. We also posit a (countably) infinite sequence of probability density functions, one for each hypothetical player: $\{p_1(\sigma), p_2(\sigma), \ldots, p_n(\sigma), \ldots\}$ where n indexes a player's relative skill, low numbers equate to low skill, higher numbers equate to higher skill. If the probability density functions $p_n(\sigma)$ embodies *all* relevant game knowledge possessed by a player of skill n, then the following limit is approached monotonically:

$$
\lim_{n \to \infty} D_s(p_n(\sigma)\,\|\,p_\infty(\sigma)) = 0 \qquad (5)
$$

That is to say the sum total of knowledge a player has embodied in $p_n(\sigma)$ which can be observed in a player's choice of moves should approach that of the oracle as skill increases. In practice there is no infinite sequence of players or player ranks from which to construct this sequence of distributions. Equally, there is no realistic single probability distribution which encompasses all game knowledge, nor is there a true oracle. Instead we group together various players into clusters based on their rank and use the distributions over moves over a variety of different board positions as a proxy for the game knowledge obtained by players at the aggregate level. Finally we take the very best players in the world as a proxy for the oracle.

IV. METHODOLOGY

The source of games we used for our analysis come from the GoGod[6] commercial database and Pandanet[7], an internet Go server for playing online. The GoGod collection has approximately 54,000 professional games and we were able to collect approximately 25,000 amateur games where the rank is 2 kyu or greater.

We collected the games into the following categories based on player ranks with the total number of games in parentheses: 2 kyu (6,262), 1 kyu (5,982), 1-2 Dan amateur (18,577), 3-8 Dan amateur (7,521), 1-4 Dan professional (4,352), 5-7 Dan professional (6,221), 8-9 Dan professional (15,761).

We considered four different, commonly occurring, patterns in Go for the analysis. The subset of the board they were taken from ranged in size from 3×5 to 9×9, i.e. a maximum of almost a quarter of the board, and these sub-regions of the board correspond to the r_i discussed earlier. We began by choosing the size of the region of the board to consider and the pattern of stones we wanted in place in that region before starting. Potentially, all of these patterns could occur at any stage of the game. However the corner positions are more likely during the beginning of the game and the smaller positions are more likely up until the end game when the board becomes quite crowded. This enables us to study patterns of play that are spread across a large range of the game. This is particularly important as the very stages of Go for which no formal results are known (e.g. the work of Berlekamp [15]) are the subject of this study.

The smallest two patterns are the Skip One pattern and the Knight pattern so named for the work in [16]. Skip One is where two stones of the same colour are separated by an empty position. The Knight configuration is two stones of the same colour separated by two empty positions in one direction and one empty position in the orthogonal direction. In Figure 1 these two starting configurations are given by the two un-numbered black stones. Note that unlike the other two starting positions (discussed next), these two positions can occur anywhere on the board.

Similarly we considered two intermediate board sizes as shown in Figure 2. These corner positions are known as Joseki, they are well studied patterns of moves and there are

many texts analysing their strategic importance. Note that the beginning configurations of these board patterns only contain two stones. These two stones in the these two patterns are in the same spatial relationship to each other as the two stones used in the smaller two pattern configurations except for the difference in stone colour, region size and the location of the board edge. These differences provide for a different context for similar physical configurations of the stones.

For each pattern of stones on the board and for each group of games taken from a set of players of a given range of ranks we constructed probability distributions of where within each region of the board the next moves were played. These distributions were then compared using the Jensen-Shannon divergence of equation 4. The most probable next move as made by the 8-9 Dan professionals is then used to update the board for all ranks. The next move distribution is derived from this new board configuration for all of the player ranks and these distributions are again compared. This process is repeated for 6 moves and the average of equation 4 is then taken over these 6 moves. This gives us the average relative entropy for each rank of players for each board configuration relative to the 8-9 Dan professionals.
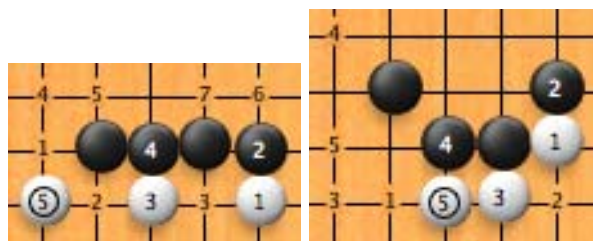

Fig. 1. The 'Skip One' continuation (left) and the 'Knight' continuation (right). The numbers are the ranked order of next moves for professionals of rank 8 or 9 Dan professionals.
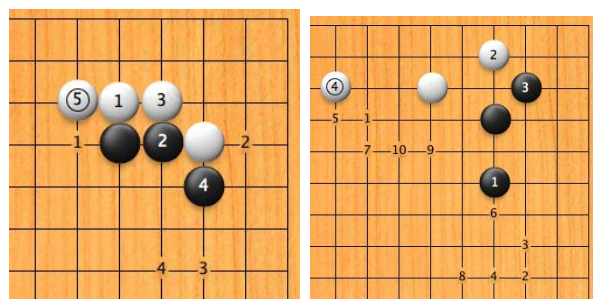

Fig. 2. The 'Avalanche' Joseki (left) and the '4-4 High Approach' Joseki (right). The numbers are the ranked order of next moves for professionals of rank 8 or 9 Dan professionals.

Looking at Figures 1 and 2 we can see both the initial configuration of the board (the stones with no numbers) as well as the most probable sequence of moves used by 8-9 Dan professionals (the numbered stones). The numbers on the board where there are no stones is the ordered set of next moves.

To demonstrate the idea, Figure 3 shows three probability distributions for where the next move will be, one for the

group of players of rank 8-9 Dan professional, one for the group of rank 1-4 Dan professional and one for the group of players of rank 3-8 Dan amateur. We label the lower ranked distributions $p(\sigma)$ and the 8-9Dan professional distribution $m(\sigma)$ and use equation 4 to measure the relative divergence of one distribution of moves form another. The 8-9 Dan professionals are the best players and we will always use them as our reference distribution $m(\sigma)$.
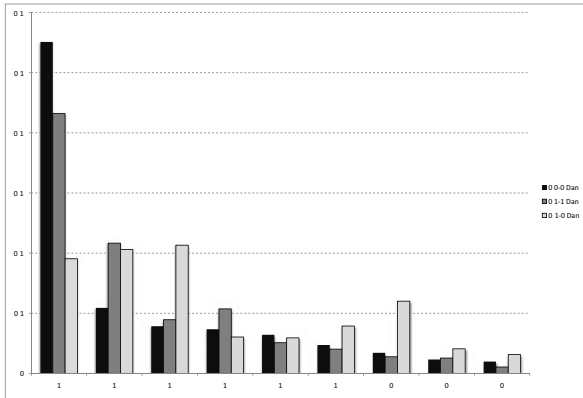

Fig. 3. Three different probability distributions for the 1^{st} move in the Avalanche Joseki. Equation 4 compares 'P 8-9 Dan' to 'A 3-8 Dan' and 'P 8-9 Dan' to 'P 1-4 Dan'
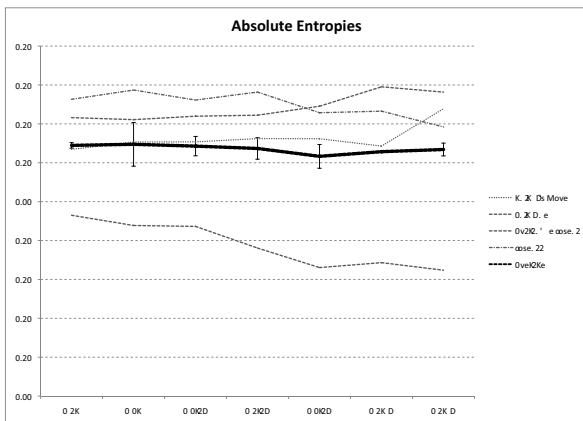

Fig. 4. A plot of the absolute entropies for all four pattern sizes, averaged over each sequence of moves. Note the 'Average' curve, taken over the other four curves, is flat to within the standard error.

V. RESULTS

Figure 4 shows the first measurements we took of the average of the absolute entropies of the distributions calculated using equation 1. The error bars are constructed from the standard error analysis for entropies as outlined in [17]. This allows us to see that there is almost no relative change in the total information content of the moves made in the various sub-regions of the boards as the rank of the players increases. Note that each curve (except the 'Average') represents an average entropy measured over a total of five or six move distributions,

so the curves in this figure aggregate 23 measures of entropy for the distribution over next moves.

Figure 5 shows the principal result of our methodology. The most important point is that the average over all board configurations and sizes (the 'Average' curve) shows a monotonic downward trend as the rank of the players approaches the 8-9 Dan professionals. This indicates that the distributions over move choices consistently moves towards that of the best players as the lower ranked players improve in their expertise. This is in line with the discussion regarding how the measure of the information contained in players move choices should approach a theoretical oracle, here taken to be the 8-9 Dan professionals.
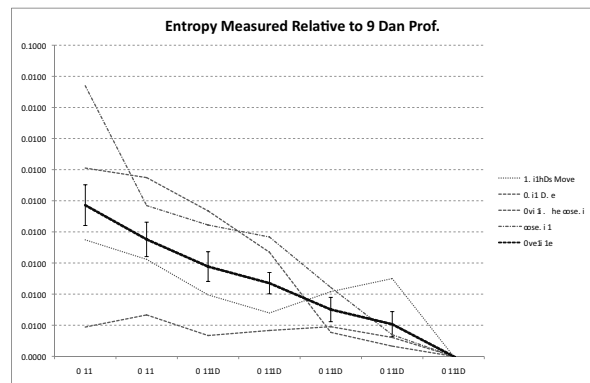

Fig. 5. The relative entropy curves and the average relative entropy curve for the four different pattern sizes used.

Figure 3 illustrates what is occurring here in the case of three different distributions over move choices taken over the same board configuration. The 3-8 Dan amateur players have an almost flat distribution across moves 1, 2 and 3. However the 1-4 Dan professionals make a very sharp distinction in their choices between move 1 and all other possible moves and this distinction is even more pronounced for the 8-9 Dan professional players. Clearly in going from amateur to professional there is a significant shift in the way players choose their moves and this change in strategic behaviour can be sharply distinguished by the relative entropy.

VI. DISCUSSION

We introduced a novel approach to the analysis of the changing strategic behaviour of players in the game of Go. Our findings show that using relative entropy reflects how the knowledge acquisition of the players, as observed in their choices over possible moves, progressively changes to more accurately correspond to the choice made by the very best players. This is an important principle which often underlies many models of knowledge acquisition used in artificial intelligence studies.

Note that these results are dependent on having taken the averages over the results we have found. As such we do not mean that specific instances of patterns or specific distributions of moves will not give different results, but that on average, over a sufficiently large number of game records and moves, our conclusions are consistent with the data.

A. The Strategic Information

However, we believe that a result of greater theoretical and practical importance can be derived from the relationship between the decreasing relative entropy of figure 5, which while interesting is also to be expected, and the flat nature of the absolute entropy. We argue that this is related to the finite capacity of a player's strategic 'notepad', i.e. the total volume of information that a player is able to manipulate when considering potential strategies.

Assume instead two alternatives to this result, that either the absolute entropy increases or decreases with player rank. Decreasing (average) entropy with player rank would imply that a skilled player is becoming (on average) more predictable as their skill increases, a counter-intuitive result given how we usually think of expertise, particularly in a competitive game. However an increasing entropy would imply that a player is becoming more random in their move choices. While there may be ways to explain such a result[8] this seems to us not very convincing.

So we look back and consider equation 5 again. We apply a further constraint on the nature of $p_n(\sigma)$: $-\sum_{\sigma_j \in \Omega} p_n(\sigma_j) \ln(p_n(\sigma_j)) = c$ for some constant c (as shown by figure 4, $c \approx 1.3$ nats or approximately 1.88 bits). With this constraint we restrict the total amount of information contained within the distribution $p_n(\sigma)$ to c and we regard c as the strategic information of our theoretical players discussed in conjunction with equation 5. With this constraint, $p_\infty(\sigma)$ of the oracle is then a perfectly optimised configuration given a finite information capacity.

Given that within our results we observe a change in player's strategic behaviour as their rank increases, and that it changes to more closely approximate the best players, an explanation is needed as to why the absolute entropy is relatively unchanged. A plausible conclusion is that our data for the relative entropy reflects equation 5 and our result for the absolute entropy reflects a finite strategic capacity of the players, i.e. expertise is not simply about acquiring and storing information, but it is tightly tied to how that information is structured and utilised within a system of limited capacity.

We are cautious regarding our interpretation of the absolute entropy. While it is clear that the average is reasonably flat, it might be the case that there are other move distributions over other board configurations which are not flat. We believe that this might well be the case, but if the entropy were calculated for each distribution for *all* possible board configurations then theoretically this encompasses all game knowledge at a certain skill level. In this way what we mean by strategic information is simply the average (or perhaps total) capacity over all configurations.

B. Constraint Development as Players Learn

We have also observed that players learn to differentiate between different strategic choices and the relationships these choices have to the board at large. As an example consider Figure 3. Here the more junior players are in a certain sense

---

[8]i.e. a skilled player is able to obfuscate their thinking to make it look random when in fact a careful strategy is being played

---

indifferent to the first 3 most likely strategies as seen by their relatively flat distribution over these 3 moves. By comparison the professionals are able to sharply distinguish between these three choices, and the 8-9 Dan professionals much more so than the 1-4 Dan professionals.

This suggests that there is at least one implicit constraint which has been learned by the professionals and enables them to strategically distinguish between these three choices. Recall that the Avalanche Joseki is contained by a corner area of size 7×7 which defines the r_i discussed earlier. The learned constraint(s) might be that either professionals understand something about the stones in the region we are considering or that they understand something about the region outside of the area considered (any part of the board other than r_i).

We consider the first of these two alternatives highly unlikely. The first move in the Avalanche Joseki is made when there are only two stones, one black one white, in r_i (see the left panel of Figure 2). It seems unlikely that the very best amateurs are unaware of the best move to play in a given region when there are only two stones present.

This suggests the reason for the professionals to differ so significantly from the amateurs is due to the professionals' more refined sensibility regarding the larger board strategy. In particular it seems likely that professionals both choose their moves within r_i and how they build other groups on the board to reflect a more directed goal for the whole board. This seems plausible as it is part of the received wisdom of Go that this is one of the hardest and last of the skills a Go player masters.

These findings are important for the purposes of machine learning in complex environments. The way in which information is integrated as well as what information is integrated is intimately tied to what we learn and when we learn it. Being able to objectively measure strategic behaviour for players of different ranks, from the very local up to a quarter of the board, will be vital to how we model learning in artificial systems.

It is also interesting to note the average relative entropy shown in Figure 5. The larger the divergence between two distributions, the larger the 'informational gap' that exists between two distributions. In practical terms this is an indicator of how much more information is required in order to move from one distribution to the other. So the smallest patterns have the smallest informational gap and the largest board patterns have the largest informational gap. This might be expected a-priori, however it also might indicate a necessary condition on how the game is learned: from the bottom up rather than from the top down.

C. Information Theory and Human Cognition

The information theoretical approach to data on player's decisions has deep relevance to human cognition. The view of human cognition and decision making as "information processing" goes back at least to the psychological studies of the 1960's where Rapaport, Messick and Fitts published a considerable body of work. For example Messick and Rapaport [18] showed a relationship between the (absolute) entropy and reward for a 10 alternative choice task. The application of information theory to the theoretical and practical study

of neural network dynamics is also an important part of the current literature (e.g. [19]).

More than this previous work, the method employed here is qualitatively different from these studies: here we have taken a real world problem where the complexity of the task tests the abilities of even the most well trained humans and we have been able to analyse their collective behaviour according to the (relative) quality of play. What we have been able to show is that the information capacity of players, at least for the common patterns analysed above, is constant across many moves in local areas. This is in agreement with the capacity limitations that has been observed in individuals [20] and has been developed in the Bayesian analysis of conceptual chunks [21]. Importantly, these cognitive limitations of human information processing have been used to build successful game playing AI systems that reflect real human limitations [8].

An interesting observation is the relationship between the (absolute) entropy $H(p(\sigma))$ of the observed distribution over moves and the mutual information $I(\sigma; \overline{\sigma})$ between σ, the set of possible outcomes of the next move, and $\overline{\sigma}$, the set of stones on the board before the next move is played. The joint probability between the prior state of the board and the next move is $q(\sigma, \overline{\sigma})$ and the marginal distribution over the prior state of the board is $p(\overline{\sigma})$. The Mutual information is defined as:

$$I(\sigma; \overline{\sigma}) = H(p(\sigma)) - H(p(\sigma)|p(\overline{\sigma})). \quad (6)$$

The second term on the right-hand-side is the conditional entropy of $p(\sigma)$ given $p(\overline{\sigma})$. This term measures how much of the information contained in $p(\sigma)$ is not explainable by knowing $p(\overline{\sigma})$. If $H(p(\sigma)|p(\overline{\sigma})) = 0$ then $p(\sigma)$ is completely determined by $p(\overline{\sigma})$, see [11]. Equation 6 measures the independence of the next move from the current state of the board. Note that $H(p(\sigma))$ is constant across the board configurations and the *local* pattern configuration (i.e. within r_i) is held constant in each case. This implies that mutual information only increases if the next move is more tightly dependent on the current board state.

We can conclude that it is the context of the board outside of r_i and how it is coupled to the next move in a local area that differentiates the rank of the players. This provides a precise way to demonstrate the qualitative discussion of the previous section regarding the global context. While this outcome may be obvious to game players, in computational terms it is not a trivial observation that Figure 5 measures the changing amount of information being used to understand the global relationships on the board in 'nats' and how these relationships influence the next move in a local context.

D. Further Work

This work will be extended in a number of different ways. First we will consider larger board patterns, possibly even whole board configurations if sufficient data can be obtained. The difficulty with these larger board spaces lies in how quickly they become unique, making reliable statistics very difficult to generate. This will enable us to establish the consistency of our results at all scales of the game. Furthermore we

wish to establish if there are any other observable phenomena which can be uncovered and analysed using this or any other information theoretical techniques. Beyond this, developing a computational model which simulates these results in an AI would usefully shed light on the advantages and disadvantages of this as a cognitive strategy.

REFERENCES

[1] Xindi Cai and Donald Wunsch, "Computer Go: A grand challenge to AI," in *Challenges for Computational Iintelligence*, pp. 443–465. Springer Berlin, 2007.
[2] Fernand Gobet, Jean Retschitzki, and Alex de Voogt, *Moves in Mind: The Psychology of Board Games*, Psychology Press, 2004.
[3] Xiangchuan Chen, Daren Zhang, Xiaochu Zhang, Zhihao Li, Xiaomei Meng, Sheng He, and Xiaoping Hu, "A functional MRI study of high-level cognition II. The game of Go," *Cognitive Brain Research*, vol. 16, no. 1, pp. 32 – 37, 2003.
[4] Yasuomi Ouchi, Toshihiko Kanno, Etsuji Yoshikawa, Masami Futatsubashi, Hiroyuki Okada, Tatsuo Torizuka, and Mitsuo Kaneko, "Neural substrates in judgment process while playing Go: A comparison of amateurs with professionals," *Cognitive Brain Research*, vol. 23, no. 2-3, pp. 164 – 170, 2005.
[5] J. Burmeister and J. Wiles, "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*, 1995.
[6] Y. Wang and S. Gelly, "Modifications of UCT and sequence-like simulations for monte-carlo Go," in *IEEE Symposium on Computational Intelligence and Games, 2007.*, 2007, pp. 175–182.
[7] L. Chang-Shing, W. Mei-Hui, G. Chaslot, J.-B. Hoock, A. Rimmel, O. Teytaud, Shang-Rong T., Shun-Chin H., and H. Tzung-Pei, "The computational intelligence of MoGo revealed in Taiwan's computer Go tournaments," in *IEEE Transactions on Computational Intelligence and AI in Games.*, 2009, vol. 1, pp. 73–79.
[8] Fernand Gobet, Peter C.R. Lane, Steve Croker, Peter C-H. Cheng, Gary Jones, Iain Oliver, and Julian M. Pine, "Chunking mechanisms in human learning," *TRENDS in Cognitive Sciences*, vol. 5, no. 6, pp. 236 – 243, 2001.
[9] C. Shannon, "Programming a computer for playing Chess," *Philosophical Magazine*, vol. 41, no. 314, 1951.
[10] J. Robson, "The complexity of Go," in *Proceedings of IFIP Congress*, 1983, pp. 413–417.
[11] T. Cover and J. Thomas, *Elements of Information Theory, 2nd edition*, Wiley, 2006.
[12] Peter D. Grünwald and A. Philip Dawid, "Game theory, maximum entropy, minimum discrepancy and robust bayesian decision theory," *The Annals of Statistics*, vol. 32, no. 4, pp. 1367–1433, 2004.
[13] S. Kullback and R. A. Leibler, "On information and sufficiency," *The Annals of Mathematical Statistics*, vol. 22, no. 1, pp. 79–86, 1951.
[14] J. Lin, "Divergence measures based on the shannon entropy," *Information Theory, IEEE Transactions on*, vol. 37, no. 1, pp. 145–151, August 2002.
[15] E. Berlekamp and D. Wolfe, *Mathematical Go: Chilling gets the last point*, AK Peters, 1994.
[16] J. S. Reitman, "Skilled perception in Go: Deducing memory structures from inter-response time," *Cognitive Psychology*, vol. 8, pp. 336–356, 1976.
[17] Mark S. Roulston, "Estimating the errors on measured entropy and mutual information," *Physica D: Nonlinear Phenomena*, vol. 125, no. 3-4, pp. 285 – 294, 1999.
[18] D. Messick and A. Rapaport, "Expected value and response uncertainty in multiple-choice decision behavior," *Journal of Experimental Psychology*, vol. 70, no. 2, pp. 224 – 230, 1965.
[19] A. Borst and F. Theunissen, "Information theory and neural coding," *Nature Neuroscience*, vol. 2, no. 11, pp. 947 – 957, 1999.
[20] G. Halford, N. Cowan, and G. Andrews, "Separating cognitive capacity from knowledge: a new hypothesis," *TRENDS in Cognitive Science*, vol. 11, no. 6, pp. 236 – 242, 2008.
[21] G. Orbán, J. Fiser, R. Aslin, and M. Lengyel, "Bayesian learning of visual chunks by human observers," *PNAS*, vol. 105, no. 7, pp. 2745 – 2750, 2008.