

Strategic planning in the game of Go using coupled non-linear oscillators

Nicholas Wilkinson^{*1}, Terry Bossomaier, Michael Harré, Allan Snyder²

Abstract. The ancient oriental game of Go is the last of the great perfect knowledge games wherein strong human players still outperform AI players, but the gap is narrowing. Superior strategic skill is keeping humans several steps ahead even as computers catch up tactically. Taking the perspective that large scale topological perception is part of the key to strategic excellence, we present a new method for generating strategic plans in the game of Go. Treating stones in play as continuous perturbations of a dynamical system composed of coupled non-linear oscillators, we exploit transitory synchrony in the wave patterns thus generated to make topological inferences about possible territories later in the game. Our system is shown, using professional game records, to focus search on appropriate areas of the space of continuations. By combining spatial (rather than game-tree) representation with dynamical lookahead, we hope to overcome some of the limitations faced by static spatial mapping and dynamical game-tree based approaches.

1 Computational architecture

Oscillator networks have been applied to a range of other scene analysis tasks [3]. The 2D network in our model consists of 19x19 Hopf oscillators [2], each coupled to their von Neumann neighbourhood (the coupling term $F(n)$) by a constant ε , and to the perturbation matrix $F(p)$, which is determined by the Go stones on the board by a constant κ giving a dynamical system of the form;

$$\dot{x}_{ij} = \gamma(\mu - r_{ij}^2)x_{ij} - \omega y_{ij} + \varepsilon F(n_{ij}) + \kappa F(p_{ij}) \quad (1)$$

$$\dot{y}_{ij} = \gamma(\mu - r_{ij}^2)y_{ij} + \omega x_{ij} \quad (2)$$

with

$$F(n_{ij}) = x_{i+1j} + x_{i-1j} + x_{ij+1} + x_{ij-1} \quad (3)$$

$$r_{ij}^2 = x_{ij}^2 + y_{ij}^2 \quad (4)$$

μ controls the amplitude of the oscillations, ω the frequency, and γ controls the speed of recovery after perturbation. Constants are set to $\gamma = 8$, $\omega = 0.13$, $\mu = 1$, $\kappa = 0.1$, $\varepsilon = 0.08$. $F(p)$ is defined by the stones on the board. Representing the board as a matrix, all white stones contribute +1, and all black stones -1, to each of their empty 4-neighbours (in Go terms, the stone's liberties). Somewhat arbitrarily, we set ω to 0.13 here. Slightly higher or lower ω gives shorter or longer lengths to the surface waves. Large changes in ω result in qualitatively different behaviour including unwanted phase

locking and a tendency towards static equilibria. At this early stage of development we simply choose settings which produce interesting and desired behaviour by hand, accepting that these are imperfect. Edge nodes are self-coupled to produce the effect of reflection at boundaries. There is no learning, in that the system does not undergo any long term changes in structure. Transitory patterns of synchrony provide the basis for a series of topological hypotheses about territory distributions given different continuations from the current position. We call these "strategic hypotheses".




Figure 1. Two representations of a professional game position at move 40




Figure 2. Some strategies suggested by our system, in which black wins. One is edged for more appealing Go-like aesthetics

2 Methods

Next we show how our system extends the capabilities of topological mapping approaches in Go [4], [1]. Using the game position in Figure 1, we generate a range of strategic hypotheses. Figure 2a-b shows 2 (of many) different strategies suggested by our system as suitable to the existing position in this game. The trajectory of the dynamical

¹ Centre for Research in Complex Systems, School of Information Technology, Charles Sturt University, Bathurst NSW 2795, email: nwilkinson3@gmail.com

² Centre for the Mind, University of Sydney

surface explores a series of variations on a "topological theme" set by the existing stones in play. We can lend normative structure to the search space using thresholded influence maps [4] of actual continuations from professional game records. Assuming that reasonable continuations are not uniformly distributed in the space of possible topologies, we consider that the neighbourhood of the actual continuation in professional games is likely to be productive for 'good' continuations. For each step of 10 moves (from move 40 to 140) into 100 professional games, we record the mean distance over the course of the run and also the lowest distance achieved, and compare these results to random growth of territories³.

3 Results

Figure 3 a shows the results of this comparison averaged over 100 games. Both the mean distance and the lowest distance are consistently smaller for the strategy maps produced by our system than for randomised continuations. The randomiser achieves a mean level about 5 percent above chance this thanks to the 40 stones already on the board, and occasionally makes an unusually lucky guess. Testing 100,000 randomised topologies we find only 11 such lucky guesses which are closer to the target (at 40 moves ahead) than the mean distance of the systematic trajectory. This suggests that the system's trajectory is limited to a sub-space somewhere in the region of 0.01 percent of that explored by the randomiser. Figure 3 b plots distance from the target for systemic vs randomised individual continuations during a run, with a target vector taken from 40 moves into the same game we have been considering above. Looking at the variation over sequential measurements, the difference in the form of the trajectory is clear. Whereas the randomiser generates at each time step a continuation which is independent of the previous continuation, the systematic trajectory passes through a series of gradually changing continuations, "variations on a theme".

4 Discussion

We interpret our results as evidence that the trajectory of the dynamical system inhabits a relatively productive subregion of the space of continuations. The reduced region in which the system concentrates its exploration will undoubtedly cut out good strategies that exist in unexplored regions, but there is no reason to believe that human players do not suffer the same limitation when they focus their search. Assuming Chinese rules, we can count unambiguously the territories of each player in a given strategic hypothesis, and so find winning strategies. Part of the definition of a good strategy is, of course, one in which we win. The other major questions are whether, and how, the strategy is tactically achievable. This question can only be addressed by further work integrating the system with tactical modules and playing against rated opponents. Tactical search can then be limited to examining the tactical viability of a limited number of specific winning strategies.

5 Conclusion

We have presented a new method for producing strategic hypotheses in computer Go. Our approach extends earlier work on visual-spatial

³ We generate random continuations by assigning each empty intersection on the board randomly to black or white, then convolving the board 5 times and thresholding to achieve areas which are similar to territories in Go. Naive Monte Carlo simulation gives a similar result.




Figure 3. Results of comparing strategies generated by topological inference versus randomised strategies

methods in Go, enabling dynamic global exploration using topological inferences based on transitory synchrony in wave patterns. We have shown that despite constraining this exploratory process to a relatively small subregion of the space of possible topologies, the system consistently follows a trajectory closer to continuations in real professional games than the random baseline. This provides evidence that this subregion is populated by "reasonable continuations", based on the assumption that real professional continuations are members of this sub-population of reasonable continuations and that this sub-population is over represented in the subspace nearby the actual continuation.

6 Acknowledgments

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

REFERENCES

- [1] B. Bouzy, 'Mathematical morphology applied to computer go', *International Journal of Pattern Recognition and Artificial Intelligence*, **17**(2), 257–268, (2003).
- [2] L. Righetti, J. Buchli, and A. J. Ijspeert, 'Dynamic hebbian learning in adaptive frequency oscillators', *Physica D*, **216**(2), 269–281, (2006).
- [3] D. Wang, 'The time dimension for scene analysis', *IEEE Trans. Neural Networks*, **16**, 1401–1426, (2005).
- [4] A. Zobrist, 'A model of visual organisation for the game of go', *Proc. AFIPS Spring Joint Computer Conference*, -, 103–111, (1969).