Zero Sum Games As Distributed Cognitive Systems

Robert L. West (rwest@hkucc.hku.hk)
Department of Psychology, University of Hong Kong, Hong Kong, PRC.

Abstract

In the case of two individuals in a competitive situation, or "game," the game itself (i.e. the players, the rules, the equipment) can be considered to constitute a distributed cognitive system. However, the dominant model of competitive behavior is game theory (VonNeumann & Morgenstern, 1944), which has traditionally treated individuals as isolated units of cognition. By simulating game playing with neural networks, and also by using human subjects, it is demonstrated that the interaction between two players can give rise to emergent properties which are not inherent in the individual players.

Recent work in distributed cognition (e.g. Hutchins, 1994; Norman, 1993; Zhang, 1997; Zhang & Norman, 1994) has indicated that cognitive processing can take place across distributed systems composed of multiple, interacting cognitive systems. For example, navigating a large ship, such as a naval vessel, is accomplished through interactions amongst specially trained humans and specialized equipment (Hutchins, 1994). Distributive systems involving more than one agent are prototypically cooperative in nature, in that the agents involved benefit from the function of the distributed system (e.g. a ship avoids sinking). However, distributed systems may also result in situations in which some individuals benefit at a cost to others. The simplest example of this is the case of two individuals in a zero sum game (i.e. a game in which only one player can win). Games such as this can be thought of as distributed cognitive systems with the goal of choosing one player as the winner.

Although game playing clearly involves interactions between the players, it does not necessarily follow that we need to consider the distributed properties of a game in order to understand the behavior of a player. This depends on whether the functionality of the cognitive mechanism used by an individual player can be understood in isolation, or needs to be interpreted in terms of the role it plays in the distributed system. The answer to this question will depend to some degree on our assumptions concerning the game playing process. For example, game theory (VonNeumann & Morgenstern, 1944) describes how rational players should behave in a competitive situation prescribed by rules and with payoffs for certain results. However, in order to do this it is necessary to make assumptions concerning the cognitive mechanisms available to the players. One assumption that is frequently made is that players have the ability to generate random responses (i.e. to draw responses at random from a predetermined distribution). For example, the game theory solution for Paper, Rocks and Scissors (hence forth PRS) is to play randomly, 1/3 paper, 1/3 rocks, and 1/3 scissors (in PRS play: paper beats rocks, rocks beats scissors, and scissors beats paper). With this assumption in place there is nothing to be gained by viewing PRS as a distributed system because players' interactions are limited to tossing out and receiving random responses. However, the assumption of random responses is problematic for two reasons. The first is that people are normally quite bad at generating random responses (see Tune, 1964, and Wagenaar, 1972 for reviews), and the second is that when people guess what is coming next in a series they attempt to capitalize on sequential dependencies, regardless if they are present or not (e.g., Anderson, 1960; Estes, 1972; Restle, 1966; Rose & Vitz, 1966; Vitz & Todd, 1967; Ward, 1973; Ward & Li, 1988).

Given the above research, a more realistic model of PRS play would have players trying to detect each others sequential dependencies. Note that the story is now different if we consider the players in isolation or if we consider them within the context of the distributed system formed by the game. Taken in isolation, a player's strategy appears passive, limited to searching for sequential dependencies in their opponents responses. However, from the distributed perspective the situation is highly interactive as each player both drives, and is driven by, their opponent's responses (i.e. my behavior would be based on my beliefs about sequential dependencies in my opponents play, which would be driven by my opponents behavior, which in turn is driven by my behavior in a similar way) . The question is, whether this highly interactive situation can impart an alternative functional significance to a sequential detection mechanism?

The Decoy Strategy

Given an opponent who is using the strategy of searching for sequential dependencies, we can ask the game theory question of how a rational opponent should respond. Generating random responses will certainly avoid any disadvantage, but it will also fail to produce an advantage. The ideal strategy under these conditions would be to use one's own responses to lure the opponent into a predictable pattern of play which could somehow be exploited. Interestingly, this agrees well with peoples' reports of how they play games such as PRS. Aside from a minority who claim to respond randomly, most people claim to deceive their opponents by allowing them to detect biases which are, in reality, decoys drawing their opponents into a predictable pattern of play. This strategy of using ones own pattern of responses to exert control over one's opponent's responses will be referred to as the decoy strategy.

What I will endeavor to show in this paper is that the function of sequential detection mechanisms within a game situation is not to passively detect sequential dependencies, but to execute the decoy strategy. Furthermore, it will be demonstrated how the ability to do this is mediated by working memory.

Simulating the Decoy Strategy

The sequential detection mechanisms assumed to be used by human game players were modeled using two layer neural networks with one layer for input and one for output (i.e. perceptrons, Rosenblatt, 1962). The output layer consisted of three nodes, to represent paper, rocks, and scissors. The input layer consisted of a variable number of three node sets. Each set represented the previous outputs of the opponent network at a particular lag, with the three nodes in each set again representing paper, rocks, and scissors. Thus the networks could be set to "remember" any number of trials back from the current trial. To represent this the networks are be referred to in terms of how many lags back they could recall (i.e. a lag1 network can remember one trial back, and a lag2 network, two trials back). Outputs were determined by summing the weights associated with the activated connections. If two or more output nodes were equally weighted the tie was resolved through random selection. Learning was accomplished through back-propagation in which a win was rewarded by adding 1 to the activated connections leading to the node representing the winning response, and a loss was punished by subtracting 1 (ties were treated as losses). In all trials, both networks began with all weights set to zero.

The neural network mechanisms used in this study were deliberately made as simple as possible in order to keep the process as transparent as possible. Also, the use of perceptrons means that the individual networks can be treated as linear systems, an important consideration for game theorists.

Simulation Results

The effect of memory was clear, networks that could remember more always won in the long term. Figure 1 displays a representative result of a lag2 network versus a lag1 network. However, as would be expected by symmetry, when a lag2 was pitted against another lag2 network no advantage emerged.

According to the decoy strategy a player wins through controlling the opponent's responses. This strategy can be seen in the causal nature of the simulation results. The network with the higher lag factor was able to win not by passively detecting sequential dependencies but by creating them. Unlike humans who might be predisposed to generating sequential dependencies, the networks based their responses solely on each others play. Thus any tendencies for one network to be predictable were *caused* by the other network.

The Decoy Strategy in Humans

The next step was to find out if human subjects could execute the decoy strategy. To do this, human subjects played PRS against a lag1 network. There were two reasons for having them play against the network instead of against each other. First, it seemed likely that they would be approximately equal in ability which, according to the simulation results above, would produce an unremarkable outcome in the long run (i.e. a 50/50 chance of winning). Second, under these conditions the only sequential dependencies present in the computer would be ones *created* by the subject. In order to win subjects would have to both create and exploit sequential dependencies in the lag1 network.

Method

Subjects The subjects were 13 volunteers from the University of British Columbia and the University of Hong Kong.

Apparatus Subjects played against a lag1 network implemented in Visual Basic (the simulations were also done using the same program). Subjects selected their PRS outputs by using a mouse to click on three different icons. Following this they clicked on a button to reveal the computer's response. The score and the number of trials were displayed so subjects could monitor their progress.

Procedure Each subject played for approximately 20 minutes. The number of trials varied based on each subject's playing speed. All subjects played at least 250 trials (mean number of trials = 441). Subjects were instructed that the computer was programmed to play like a human, and that it was possible to beat it. They were also told that the program was very complex and that they should play by intuition.

Results

Figure 2 displays the subject's score minus the computer's score across trials. The data was combined so that each subject's game picks up where the previous subject left off (e.g. subject 1 finished with a lead of 44 points after 800 trials so subject 2 was plotted as though he began with a 44 point lead starting at trial number 801). This was done in order to get a sufficient number of trials (total number of trials = 5727) to indicate an unambiguous trend. The upward trend in Figure 2 is very clear and demonstrates that the human subjects were able to execute the decoy strategy.

Discussion

The neural networks used in this study were designed to passively detect sequential dependencies. The decoy strategy was not implicit in the design of these networks, but emerged from the *interaction* between them. Although it is possible that the human subjects were able to win by some other means it is unclear how this could be achieved. Also, it is doubtful if other explanations could achieve the same level of parsimony, or consistency with previous research.

The implications of these results go far beyond describing a good strategy for playing PRS, as it is possible that a considerable amount of competitive behavior is based on this type of process. More generally, the results of this study are consistent with the view that human cognition needs to be understood within in the environmental context

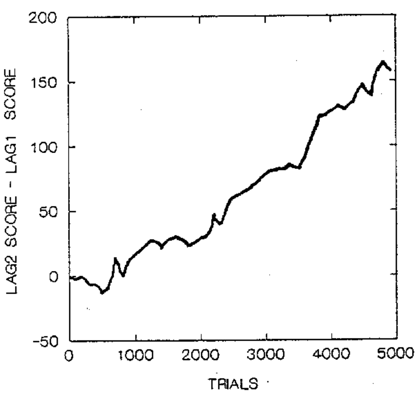

Figure 1: A Lag2 network versus Lag1 network

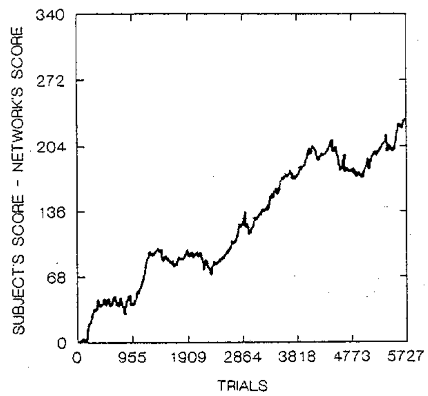

Figure 2: Human subjects versus a Lag1 network

in which it developed (e.g. Gibson, 1986). For humans this entails understanding an individual within a social context that is both cooperative and competitive. As demonstrated in this study, the benefit of an individual cognitive system may reside in the type of distributed system it creates when joined with other systems, rather than in its function as an isolated unit.

The advantage of the methodology used in this study is that it reconciles distributive and individual cognitive research for this type of behavior. Using the detailed findings of traditional cognitive research on individuals, tentative models can be constructed and placed in interactive situations. The emergent patterns from such simulations can then be compared to simulations in which one of the simulated agents is replaced with a human agent, or to interactions between two humans. In this way we can begin to understand the relationship between individual cognitive agents and the emergent, distributed systems in which we live.

References

Anderson, N. H. (1960). Effect of first-order probability in a two choice learning situation, *Journal of Experimental Psychology, 59*, 73-93.

Estes, W. K. (1972). Research and theory on the learning of probabilities. *Journal of the American Statistical Association, 67*, 81-102.

Gibson, J. J. (1986). *The ecological approach to visual perception*. Hillsdale, NJ: Erlbaum.

Hutchins, E. (1994). *Cognition in the wild*. Cambridge, MA: The MIT Press.

Norman, D. A. (1993). *Things that make us smart*. Reading MA: Addison-Wesley.

Restle, F. (1966). Run structure and probability learning: Disproof of Restle's model. *Journal of Experimental Psychology, 72,* 382-389.

Rose, R. M., & Vitz, P. C. (1966). The role of runs of events in probability learning. *Journal of Experimental Psychology, 72*, 751-760.

Rosenblatt, F. (1962). *Principles of neurodynamics*. New York: Spartan.

Tune, G. S. (1964). A brief survey of variables that influence random generation. *Perception and Motor Skills, 18*, 705-710.

Vitz, P. C., & Todd, T. C. (1967). A model of learning for simple repeating binary patterns. *Journal of Experimental Psychology, 75*, 108-117.

VonNeumann, J., & Morgenstern, O. (1944) *Theory of games and economic behaviour*. Princton, N. J.: Princton University Press.

Wagenaar, W. A. (1972). Generation of random sequences by human subjects: A critical survey of the literature. *Psychological Bulletin, 77*, 65-72.

Ward, L. M. (1973). Use of markov-encoded sequential information in numerical signal detection. *Perception and Psychophysics, 14,* 337-342.

Ward, L. M., Livingston, J. W., & Li, J. (1988). On probabilistic categorization: The markovian observer. *Perception and Psychophysics, 43,* 125-136.

Zhang, J. (1997). A distributed representation approach to group problem solving. *Journal of the American Society of Information Science*, in press.

Zhang J. & Norman, D. A. (1994). Representation in distributed cognitive tasks. *Cognitive Science, 18 (1)*, 87-122.