

Evolvable Poker Agent

Ciarán Kelly^{*}
Artificial Life Laboratory
RINCE Group
Dublin City University
Glasnevin
Dublin 9
ciaran.kelly9@mail.dcu.ie

ABSTRACT

The popular poker variant, Texas Hold Em', can be played by 2 or more agents. Each agent is dealt 2 cards face-down, and, during 3 subsequent betting rounds, a further 5 "community" cards are dealt face-up(3-1-1) for the use of all players. The objective is to form the best hand using any combination of the 2 so-called "hole" cards, and the 5 community cards. As with all poker variants, the winning hand does not necessarily constitute the **best hand**, in the situation where the other players fold in the face of betting/raising from another player, who might, of course, be bluffing. At each betting round, players may choose from the following actions: Check, Bet, Raise, Fold. (Checking is an action which passes play to the next player, but costs no money to do so. Checking is only possible when no other player has so far bet.) In terms of Classical AI, the games of Chess and Draughts are seen as games of perfect information, where the entire state space could theoretically (even though not practically) be searched to analyse every possible move. Poker, on the other hand is a game of imperfect information, as we have no way of knowing which cards other players have been dealt, and are forced to base our plays on both the strength of our own cards, and the perceived strength of the opponents cards. Seen by many as a game of skill, Texas Hold Em' is sometimes seen as one of those games where you can beat the house, unlike Blackjack or Roulette, which rely almost entirely on luck (even **with** card counting!)

1. Introduction

It is not surprising that there are not many good poker playing AI agents floating around. The few [?][?] that do exist are either specialised and optimised for a particular type of Hold Em' play— eg. no-limit head to head— or tend to base themselves heavily on mathematical algorithms which depend on only a few variables of the game. I would like to

^{*}Student Number: 99290146

see this poker agent idea taken down the ALife road as a complement to the AI road. It should be possible to evolve a population of poker playing agents by simulating a tournament type playing structure, with the winning players advancing, reproducing, whatever, while the losing programs (ie. the ones that have played themselves into bankruptcy!) effectively dying off.

2. Applicability of ALife to the Problem

I am a firm believer that the environment has a strong part to play in the evolution of living agents, to the point that evolution will tend towards the "goal" of the environment. By "goal", I mean to say whatever survival criteria we have set for the agents. If we have an environment that favours agents of small code length, then small code length agents will proliferate. The same could be said for any environment which has as it's central "goal" one single characteristic. On the other hand, a game such as poker is a much more realistic environment to attempt evolution in. Ray's Tierra [?] used the concept of the "Reaper" to kill off agents that had reached old age. While performing the essential task of clearing out space for new agents, the "Reaper" is blind to the relative strengths of agents, to a point (Agents, through reproduction, can stave off the advances of the Reaper, but ultimately are destined for a final appointment with it.) It could be argued that software agents "live" in an environment where time passes in a fashion greatly different to our own, relatively speaking of course. In the circumstances of a game of poker, evolved agents could theoretically evolve the ability to bluff, call bluffs and so on, whereas their AI counterparts may be stifled by these tasks.

3. Evolving a Player

At it's most basic level (assuming we already have an artificial environment capable of storing the agents and facilitating multiple games of poker), I envisage the agents participating in a game somewhat akin to the Prisoners Dilemma [?], with each player folding in turn, as this is possibly the best group strategy for staying alive (the starting position rotates so each round, a different player will be left standing and take the pot). At some stage, one or more of the players may evolve to take an action other than fold, thus ensuring they win the pot each time. These non-folders would pass their genes down the line, and presumably the non-folders would displace the folders. The idea of hand-strength is clearly not mentioned here and I believe that

agents would need to evaluate how strong they believe themselves to be (Afterall, the animal kingdom displays many examples of creatures displaying their dominance over one another). After assessing their own strength, the agents could decide what actions to take. This points out another oversight, if you'd like to call it that. How do agents select their actions? Neural Nets would seem a logical solution, and I feel that they may very well offer the best solution. If what they say about Neural Nets is true, neural nets can simulate a basic collection of brain cells. A method of representing a neural net in a "genome" will need to be designed to facilitate the reproduction process along with all that goes with it (cross-over, slight possibility of random mutation, etc).

4. The Circle of Life

A system such as the one described above could also allow for "mate selection", with programs trying to mate with the best of the other programs available. In this respect, programs would not be explicitly told to look for better programs, they would need to evolve that capability. Some sort of data structure would be available within each program to display its characteristics for others to see (such as chip total, success rates at different stages of the game, games played (age), number of offspring). If a program were to wear its colours on its sleeve so to speak, prospective mates may evolve to choose stronger programs to mate with. At this point, we should be careful not to impose any requirements on **how** agents choose their mates, if they want to choose the least successful, so be it, we are more interested in emerging patterns than **hand designing** the "best" agent. This display of characteristics may also give rise to some sort of dominance relationship between agents, with some agents being almost "afraid" of others, or vice-versa, some agents take bigger risks in the face of less successful agents.

5. Conclusion

It is important to remember that this would be an experiment in ALife rather than AI. The game of Texas Hold Em' poker has been chosen because it is relatively easy to learn, yet conceals some complex strategies which can be picked up while playing. The game also provides a more complicated environment for the agents to interact with and hopefully will kickstart them into evolving some interesting properties. You will remember that for these agents, survival is a constant struggle against other agents. Survival is not only about reproduction. Clearly there is a "survival of the fittest" element in the immediate sense, but also in the long-term evolutionary sense, in that only the strongest agents will survive to reproduce, on average.

The goal of this project would not necessarily require the evolution of an agent of comparable skill to a human. Certainly, if evolution in a system like this were to take off, it could be imagined that a skilled agent or an entire species of skilled agents could emerge, but for some reasons in my opinion, this may not actually happen. The agents are, afterall, only in competition with themselves, and not humans, and therefore could not be expected to evolve strong strategies against humans. On the other hand, given a complex and realistic enough environment, I believe that the agents could evolve to just **be good** at poker, regardless of opponent.

The complexity of evolving such a "brain" becomes clear

when one considers that all that is passed down from generation to generation is the configuration and layout of the neural-cells. Knowledge acquired during the lifetime of an agent cannot be passed by genetic means to its offspring, unless of course you believe in Lamarckism[?] over Darwinism[?]. In the natural world, parents may teach their offspring to carry out certain tasks like hunting, nest building etc, or children may learn from their peers, in a manner akin to "play-fighting" lion cubs. In other cases, the offspring may be left totally alone to learn to fend for themselves. Although these concepts appear to be a very slow way of evolving something, they **are** the natural way, and therefore require some consideration in the model. Perhaps "newborns" could attend a PokerCreche for a set number of games, allowing them to learn before they go out into the world. The PokerCreche could also evolve, with the parent agents having an input of sorts into the way the PokerCreche is run. Theoretically, the system could survive without the creche, but I think that evolution may take longer to reach optimum configurations.