

Evolutionary Computation for Poker AI



Abstract

Our goal in this project was to teach a computer how to play Texas hold 'em poker using principles of artificial intelligence. The plan was to give the computer the rules of poker and, by using artificial neural networks and genetic algorithms, have the computer teach itself how to play poker at a respectable level. So far as possible, we avoided using huge amounts of domain knowledge (human knowledge on how to play poker well). Through many generations, storing and reproducing the strong genes and discarding the weak genes, our aim was to evolve better and smarter poker players. Varying and fine-tuning our techniques allowed us to find what worked best.

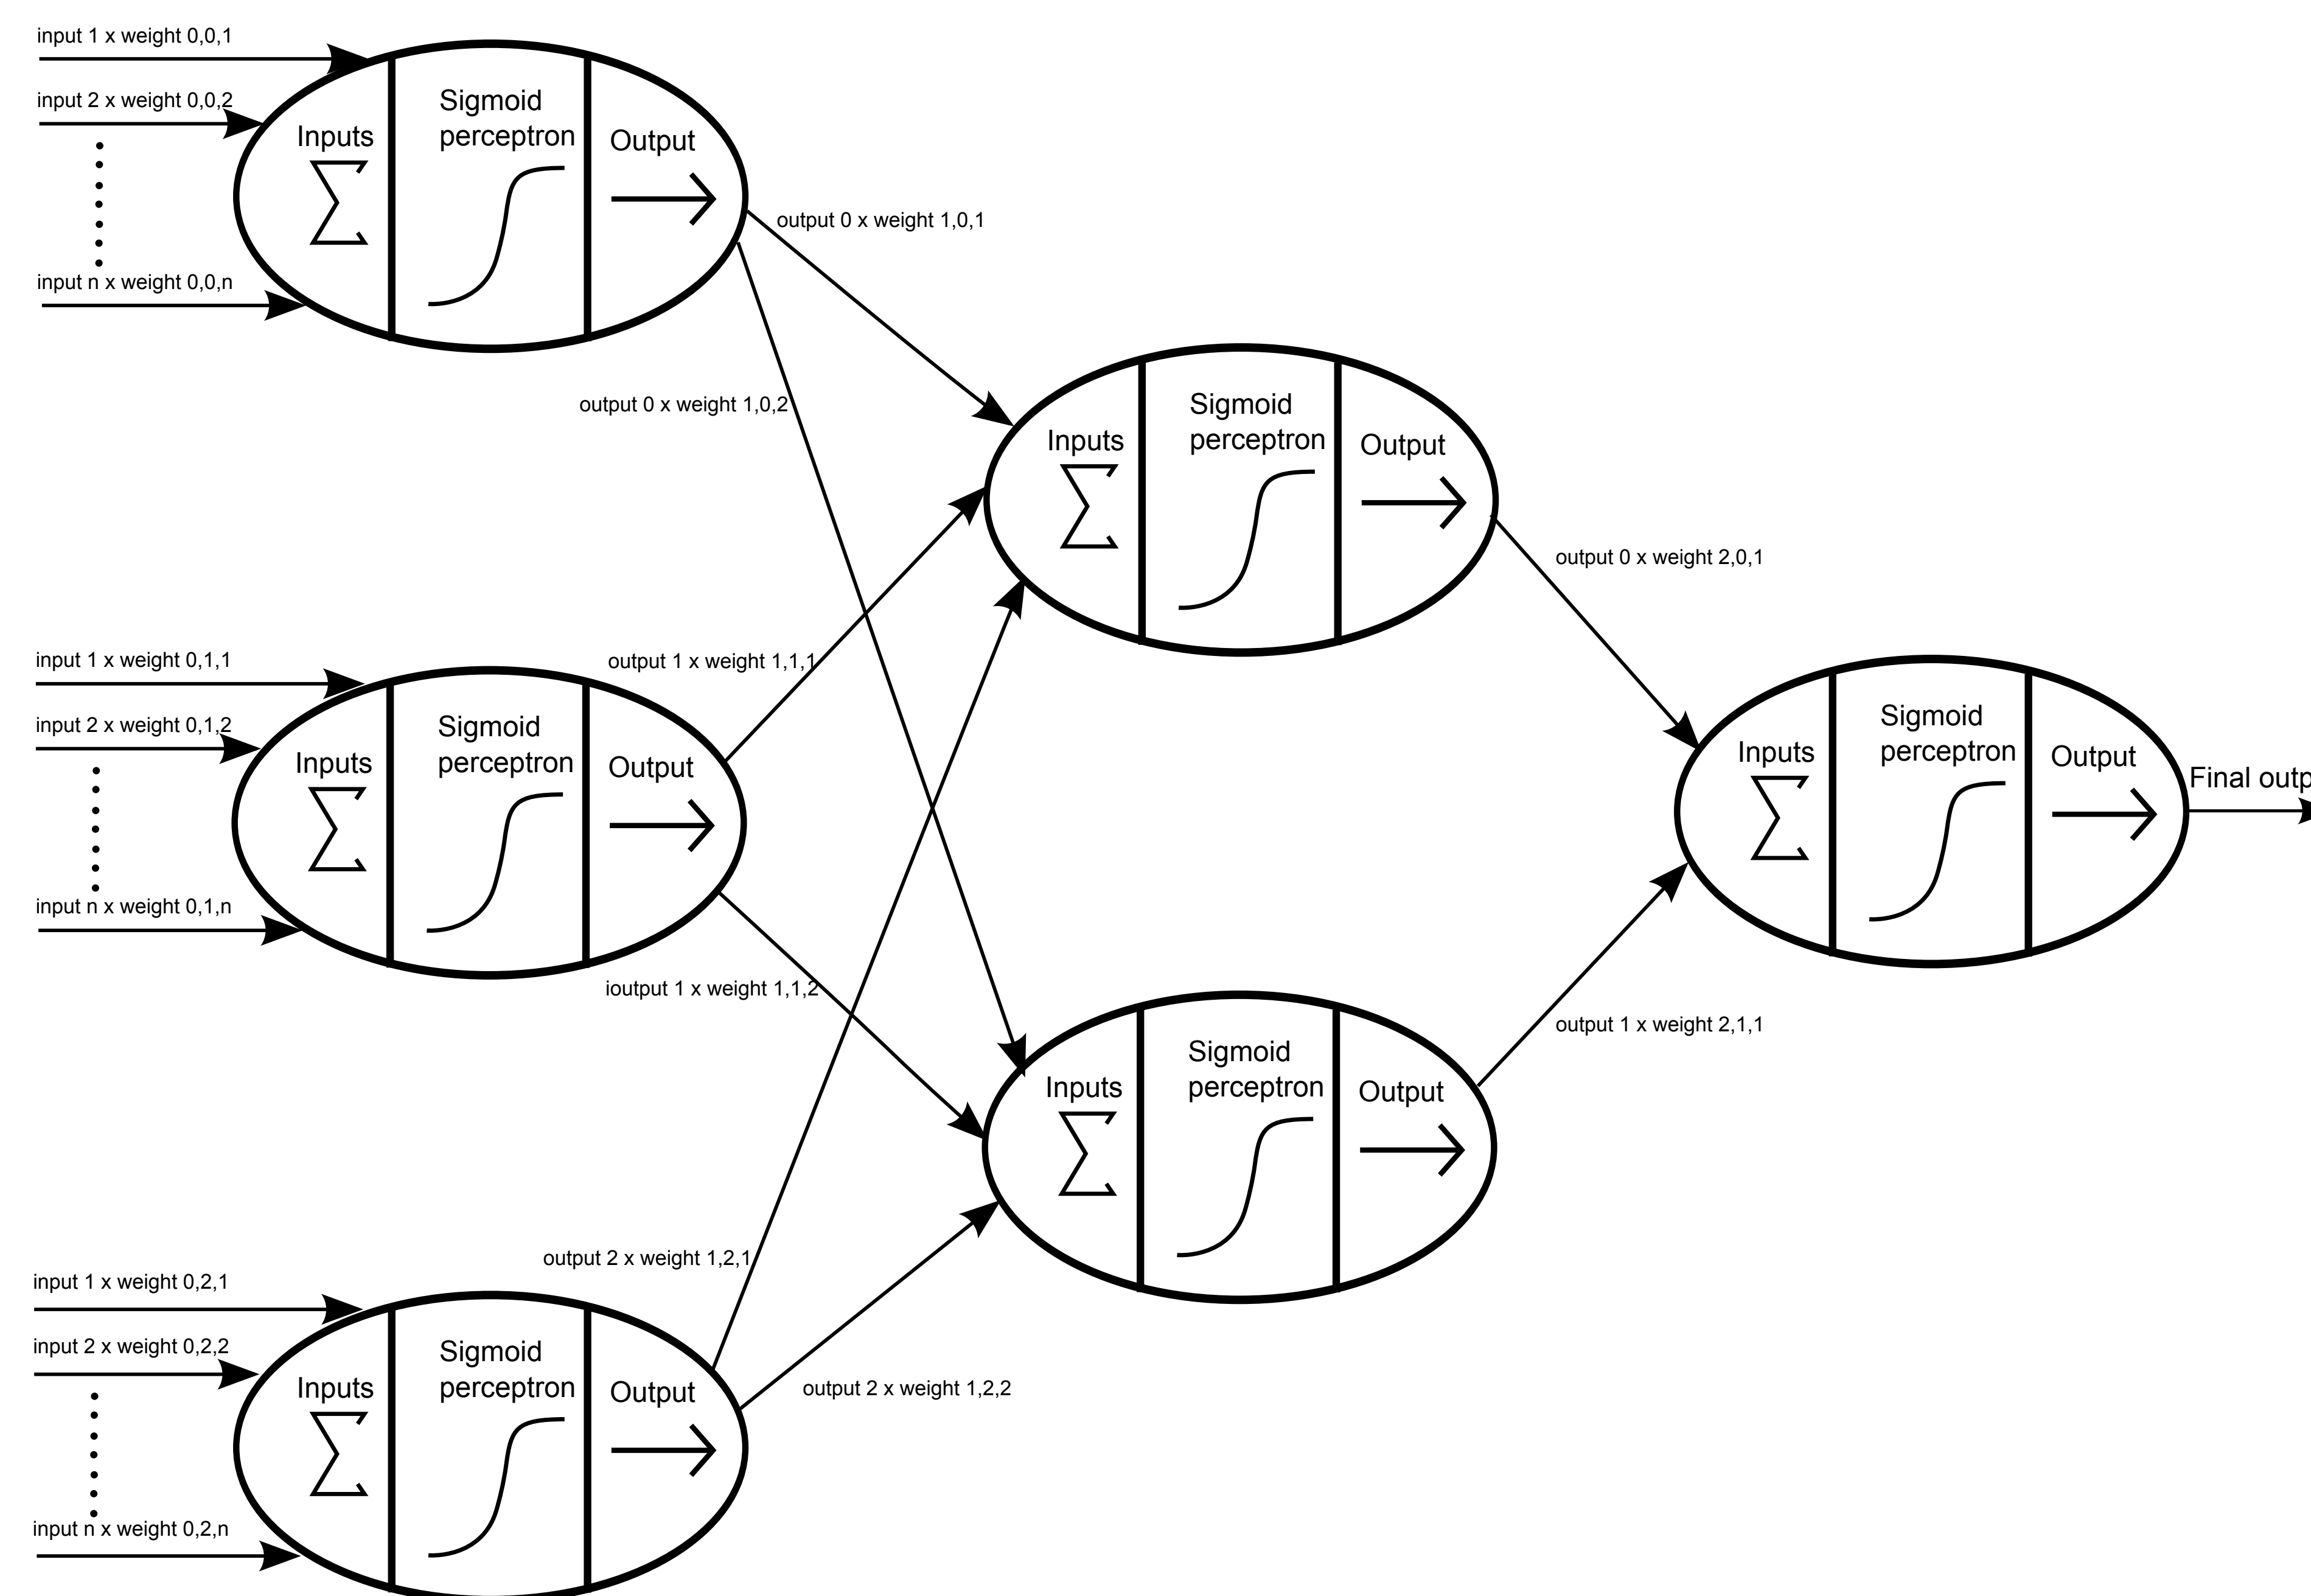
Poker Environment

The game of poker will be played in the style of Texas hold 'em. There are five players at each game, and each game continues until there is only one player left.

Each player starts with \$1,000 and is eliminated when he is out of money. There are no limits to the bets.


Artificial Neural Network

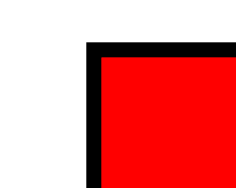

The neural network is what makes the decisions for the poker playing agent. How much the agent is willing to bet is computed by feeding inputs through a system of connected neurons. Information about the poker game like: cards, money, amount of players in the hand, are being used as inputs for the first layer of neurons in the network. Each input is multiplied with a connected weight before they are added together and passed through a sigmoid perceptron ($\frac{1}{1 + e^{-x}}$). The output of this sigmoid perceptron, always between 0 and 1, is being used as input for neurons in the next layer. The different weights in the neural network are what determine how the agent is going to play.




Genetic Algorithms

The genetic algorithm is what allows the agents to evolve and become better poker players over time. Each agent is created with a random set of genes (in our case the weights in the neural network). Through playing poker, we would discover, store, and reproduce strong genes, and discard the weak genes. The reproduction would consist of taking genes from two strong agents, and combining them to create a new agent, while the weak agents would be discarded. This works much like survival of the fittest. There is also a slight chance for gene mutation to broaden our search for strong genes. The reproduction of genes and the games are local. 100 players (10 x 10 grid) all share their turn to host a game. They invite their neighbors to play, and during evolution genes are only spread to neighboring agents.



 The host
 Invited neighbors



Simon Olsen
 Angelo State University
 solsen2@angelo.edu

