Master's Seminar banner

"Mastermind by Importance Sampling and Metropolis-Hastings"

Angela Snyder
Department of Statistics, University of Chicago

FRIDAY, October 8, 2004, 11:00 am
Eckhart Hall, Room 110, 5734 S. University Avenue

ABSTRACT

Two machine learning algorithms are introduced for playing a one player version of the logical game Mastermind. In the classic game of Mastermind there are two players: the encoder and the decoder. The encoder builds a secret code by selecting a sequence of four pegs, each chosen from six different colors. The decoder then attempts to guess the secret code in as few turns as possible. In this version, the secret code is generated uniformly at random from the set of all possible codes and the computer must guess this code in order to win the game.

The first strategy presented makes guesses according to an exponential distribution on a simple set of features. The model parameters were trained by performing a gradient descent algorithm on an importance sampling estimate of the mean number of turns required to win the game. The second strategy is based on the Metropolis-Hastings algorithm, and uses a simple cost function to evaluate each code proposed. This strategy requires more turns on average to complete the game, but provides a significant reduction in the number of codes that must be evaluated at each turn. This reduction in computation allows for fast play on versions of the game with more pegs and more colors.

Lastly, the performance of the Metropolis-Hastings based strategy is examined on a modified version of the game where false information is given with some positive probability, and the performance is compared to that of human subjects.