

An Introduction to Monte Carlo Techniques in Artificial Intelligence - Part I


Todd W. Neller
Gettysburg College

Monte Carlo (MC) Techniques in AI

- **General: Monte Carlo simulation for probabilistic estimation**
- **Machine Learning: Monte Carlo reinforcement learning**
- **Uncertain Reasoning: Bayesian network reasoning with the Markov Chain Monte Carlo method**
- **Robotics: Monte Carlo localization**
- **Search: Monte Carlo tree search**
- **Game Theory: Monte Carlo regret-based techniques**

Monte Carlo Simulation

- Repeated sampling of stochastic simulations to estimate system properties
- Recommended Readings:
 - Wikipedia article on [Monte Carlo Methods](#)
 - [Paul J. Nahin's [Digital Dice: Computational Solutions to Practical Probability Problems](#) is a great source of MC simulation exercises.]



Why MC Simulation?

- Nihin's motivational philosophical theme:
 1. No matter how smart you are, there will always be probabilistic problems that are too hard for you to solve analytically.
 2. Despite (1), if you know a good scientific programming language that incorporates a random number generator (and if it is good it will), you may still be able to get numerical answers to those "too hard" problems.

Problem Solving Approach

1. Program a single simulation with enough printed output to convince you of the correctness of your model.
2. Add your statistical measure of interest and test its correctness as well.
3. Remove printing from the code.
4. Wrap the code in a loop of many iterations.
5. Add printing to summarize the analysis of the collected statistical data.

Game AI Exercises

- Yahtzee
 - probability of getting a Yahtzee (5 of a kind) in 3 rolls of 5 dice
- Pig
 - probability of turn outcomes of “hold at 20” policy
 - expected number of turns in solitaire play
 - first player advantage assuming “hold at 20” policy
- Risk
 - attack rollouts with varying attackers, defenders
- Limitations of MC Simulation
 - probability of rolling all 1s for n dice

MC Reinforcement Learning

- Learn essential Reinforcement Learning (RL) terminology from a variety of sources:
 - Sutton, R.S. and Barto, A.G. [Reinforcement Learning: an introduction, Chapter 3](#)
 - Kaelbling, L.P., Littman, M.L., and Moore, A.W. [Reinforcement learning: a survey](#), sections 1 and 3.1
 - Russell, S. and Norvig, P. [Artificial Intelligence: a modern approach](#), 3rd ed., section 17.1
- Read specifically about MC RL:
 - Sutton, R.S. and Barto, A.G. [Reinforcement Learning: an introduction, Chapter 5](#)

Approach N

- Since learning is best through experience, we suggest implementing Sutton and Barto's MC RL algorithms with a single running problem.
- Approach N
 - Original design as simplest "Jeopardy approach game" [[Neller & Presser 2005](#)] prototype
 - 2 players and a single standard 6-sided die (d6).
 - Goal: approach a total of n without exceeding it.
 - 1st player rolls a die repeatedly until they either (1) "hold" with a roll sum $\leq n$, or (2) exceed n and lose.
 - 1st player holds at exactly $n \rightarrow$ immediate win
 - Otherwise 2nd player rolls to exceed the first player total without exceeding n , winning or losing accordingly.
- Only 1st player has a choice of play policy.
- For $n \geq 10$, the game is nearly fair.
- Sample solution output given for $n = 10$, but students may be assigned different n .

MC RL Approach N Exercises

- Comparative MC Simulation
 - Simulate games with 1st player holding sum s for s in $[n - 5, n]$. Which s optimizes 1st player wins?
- First-visit MC method for policy evaluation
- MC control with exploring starts (MCES)
- Epsilon-soft on-policy MC control
- Off-policy MC control

Further MC RL Game AI Exercises

- Hog Solitaire
 - Each turn, roll some chosen number of dice. Score only rolls with no 1s. How many dice should be rolled so as to minimize the expected number of turns to reach a goal score?
- Pig Solitaire
 - As above, but with individual die rolls and option to hold and score at any time.
- Yahtzee or Chance
 - Assuming an option to score a Yahtzee (5-of-a-kind, 50 pts.) or Chance (sum of dice) in 3 rolls, which dice should be rerolled in any given situation?

Conclusion

- Deep knowledge comes best from playful experience.
 - *“One must learn by doing the thing; for though you think you know it, you have no certainty, until you try.”* – Sophocles
 - *“Play is our brain’s favorite way of learning.”* – Diane Ackerman
- We have provided novel, fun Game AI exercises that
 - cover essentials in MC Simulation and MC RL
 - range from CS1-level to advanced AI exercises
 - have Java solutions available to instructors
 - suggest many starting points for undergraduate research projects