

Monte Carlo Tree Search

default policy: random


bias ployout
toward
better
moves

move-averaged sampling technique (MAST)
maintain reward for actions during ployout

PAST (predicate averaged sampling techniques)

	\tilde{r}	r
P_1 : A has more in store	100 101	0.52 $\frac{53}{101}$
P_2 : B has more in store	50	0.4
P_3 : A has no pit w/ > 2 seeds	50	0.38
P_4 : $\sim P_3$	100 101	0.55 $\frac{56}{101}$
P_5 : A has empty pits	50 51	0.7 $\frac{36}{51}$
P_6 : $\sim P_5$	100	0.3


P_1 satisfied 100, 0.52
 P_4 satisfied +100, 0.55
 P_5 satisfied +50, 0.70
250



suppose this is chosen in ployout and leads to W

stochastic games:

choices of which dice to keep



tree policy:
random at stochastic positions
UCB at other positions

imperfect information games:

determinization: pick values for unknown info





MCTS tree data \rightarrow playability of game?


variance of
score, spread

not too hard to win
not unbalanced to one player
not too long or too short

MCDS?

directed acyclic graph
(DAG)





implementation: edges are implicit



store as map pos \rightarrow statistics

GPU Programming

task parallel : different tasks (functions) operating simultaneously on data
 ↳ each with independent point of execution

GPU → data parallel : same task operating simultaneously on different partition of data
 ↳ at same point of execution

for $i=0$ to $n-1$
 $c[i] = a[i] + b[i]$



processor (i,j) computes c_{ij} by
 $tot = 0$
 for $k=1$ to n
 $tot += a_{ik} \cdot b_{kj}$
 $c_{ij} = tot$