Adversarial Search and Game Playing (Respect your opponent to make good decisions)

Russell and Norvig: Chap. 5, Sect. 5.1 - 5.4

Slides adopted from Jean-Claude Latombe at Stanford University (used with permission)

Games

Games like Chess or Go are compact settings that mimic the uncertainty of interacting with the natural world For centuries humans have used them to exert their intelligence Recently, there has been great success in building game programs that challenge human supremacy

Specific Setting

Two-player, turn-taking, deterministic, fully observable, zero-sum, time-constrained game



Search Problem Formulation

Initial state: Game setup at start Player(s): Which player moves in state Action(s): Legal moves in a state Result(s,a): Transition model Terminal-Test(s): True when game over Evaluate(s, p): Estimate of how good s is for player p

MIN Competes with MAX

MIN wants MAX to lose (and vice versa)

No plan exists that guarantees MAX's success regardless of which actions MIN executes (the same is true for MIN)

Time Limit

At each turn, the choice of which action to perform must be made within a specified time limit

The state space is enormous:
 only a tiny fraction of this space
 can be explored within the time limit





In general, the branching factor and the depth of terminal states are large

- Number of states: ~10⁴⁰
- Branching factor: ~35
- Number of total moves in a game: ~100

Choosing an Action: Basic Idea

Using the current state as the initial state, build the game tree uniformly to the maximal depth h (called horizon) feasible within the time limit

- Evaluate the states of the leaf nodes
- Back up the results from the leaves to the root and pick the best action assuming the worst from MIN
- → Minimax algorithm

Evaluation Function

◆ Function e: state s→number e(s)
 ◆ e(s) is a heuristics that estimates how favorable s is for MAX

e(s) > 0 means that s is favorable to MAX (the larger the better)
e(s) < 0 means that s is favorable to MIN
e(s) = 0 means that s is neutral

Example: Tic-tac-Toe

e(s) = number of rows, columns, and diagonals open for MAX
number of rows, columns, and diagonals open for MIN



Creating an Evaluation Function

Usually a weighted sum of "features":



Features may include

- Number of pieces of each type
- Number of possible moves
- Number of squares controlled





Why using backed-up values?

At each non-leaf node N, the backed-up value is the value of the best state that MAX can reach at depth h if MIN plays well (by the same criterion as MAX applies to itself)

If e is to be trusted in the first place, then the backed-up value is a better estimate of how favorable STATE(N) is than e(STATE(N))

Minimax Algorithm

- Expand the game tree uniformly from the current state (where it is MAX's turn to play) to depth h
- Compute evaluation function at every leaf
- Sack-up the values from the leaves to the root of the tree as follows:
 - MAX node→maximum evaluation of its successors
 - MIN node→minimum evaluation of its successors
- Select the move toward a MIN node that has the largest backed-up value

Minimax Algorithm

- Expand the game tree uniformly from the current state (where it is MAX's turn to play) to depth h
 Horizon: Needed to return a decision within allowed time
- Sack-up the values from the leaves to the root of the tree as follows:
 - MAX node maximum evaluation of its successors
 - MIN node minimum evaluation of its successors
- Select the move toward a MIN node that has the largest backed-up value

Game Playing (for MAX)

- Repeat until a terminal state is reached
- Select move using Minimax
- Execute move
- Observe MIN's move

Note that at each cycle the large game tree built to horizon h is used to select only one move

All is repeated again at the next cycle (a sub-tree of depth h-2 can be re-used)







The beta value of a MIN node is an upper bound on the final backed-up value. It can never increase









Alpha-Beta Pruning

Explore the game tree to depth h in depth-first manner

Back up alpha and beta values whenever possible

Prune branches that can't lead to changing the final decision

Alpha-Beta Algorithm

Update the alpha/beta value of the parent of a node N when the search below N has been completed or discontinued

Discontinue the search below a MAX node
 N if its alpha value is ≥ the beta value of a
 MIN ancestor of N

Discontinue the search below a MIN node
 N if its beta value is ≤ the alpha value of a
 MAX ancestor of N


































How much do we gain?

Assume a game tree of uniform branching factor b

Minimax examines O(b^h) nodes, so does alpha-beta in the worst-case

How much do we gain?

The gain for alpha-beta is maximum when:

 The MIN children of a MAX node are ordered in decreasing backed up values

 The MAX children of a MIN node are ordered in increasing backed up values

Then alpha-beta examines O(b^{h/2}) nodes [Knuth and Moore, 1975]

How much do we gain?

Sut this requires an oracle (if we knew how to order nodes perfectly, we would not need to search the game tree)

If nodes are ordered at random, then the average number of nodes examined by alpha-beta is ~O(b^{3h/4})

Heuristic Ordering of Nodes

Order the nodes below the root according to the values backed-up at the previous iteration

Other Improvements

- Adaptive horizon + iterative deepening
- Extended search: Retain k>1 best paths, instead of just one, and extend the tree at greater depth below their leaf nodes (to help dealing with the "horizon effect")
- Singular extension: If a move is obviously better than the others in a node at horizon h, then expand this node along this move
- Use transposition tables to deal with repeated states
- Null-move search

Checkers: Tinsley vs. Chinook

Name:Marion TinsleyProfession:Teach mathematicsHobby:CheckersRecord:Over 42 yearsIoses only 3 gamesof checkers

World champion for over 40 years

Mr. Tinsley suffered his 4th and 5th losses against Chinook

Chinook

First computer to become official world champion of Checkers!

Chess:

Kasparov vs. Deep Blue

5'10″	Height	6' 5"
176 lbs	Weight	2,400 lbs
34 years	Age	4 years
50 billion neurons	Computers	32 RISC processors
		256 VLSI chess engines
2 pos/sec	Speed	200,000,000 pos/sec
Extensive	Knowledge	Primitive
Electrical/chemical	Power Source	Electrical
Enormous	Ego	None
1997 · Deen Blue	wins hy 3 wins	1 loss and 2 draws

Chess: Kasparov vs. Deep Junior

Deep Junior

8 CPU, 8 GB RAM, Win 2000 2,000,000 pos/sec Available at \$100

August 2, 2003: Match ends in a 3/3 tie!

Othello: Murakami vs. Logistello

Takeshi Murakami World Othello Champion

1997: The Logistello software crushed Murakami by 6 games to 0

Go: Goemate vs. ??

Name: Chen Zhixing Profession: Retired Computer skills: self-taught programmer Author of Goemate (arguably the best Go program available today)

Gave Goemate a 9 stone handicap and still easily beat the program, thereby winning \$15,000

Jonathan Schaeffer

Go: Goemate vs. ??

Name: Chen Zhixing Profession: Retired Computer skills:

Go has too high a branching factor for existing search techniques

est

Current and future software must rely on huge databases and pattern-recognition techniques

handicap and still easily beat the program, thereby winning \$15,000

Jonathan Schaeffer

Secrets

Many game programs are based on alphabeta + iterative deepening + extended/singular search + transposition tables + huge databases + ...

For instance, Chinook searched all checkers configurations with 8 pieces or less and created an endgame database of 444 billion board configurations

Secrets

The methods are general, but their implementation is dramatically improved by many specifically tuned-up enhancements (e.g., the evaluation functions) like an F1 racing car

Perspective on Games: Con and Pro

Chess is the Drosophila of artificial intelligence. However, computer chess has developed much as genetics might have if the geneticists had concentrated their efforts starting in 1910 on breeding racing Drosophila. We would have some science, but mainly we would have very fast fruit flies.

Saying Deep Blue doesn't really think about chess is like saying an airplane doesn't really fly because it doesn't flap its wings.

Drew McDermott

John McCarthy

Other Types of Games

 ♦ Multi-player games, with alliances or not
♦ Games with randomness in successor function (e.g., rolling a dice) → Expectminimax algorithm
♦ Games with partially observable states (e.g., card games) → Search of belief state spaces