
COMP 102: Excursions in Computer Science

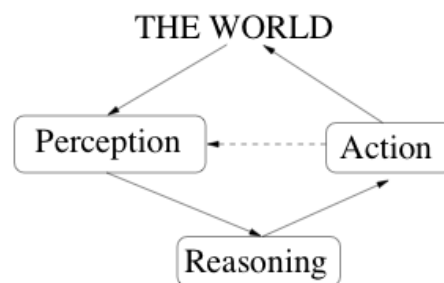
Today's topic: Algorithms for Game Playing

Instructor: Joelle Pineau (jpineau@cs.mcgill.ca)

Web page: www.cs.mcgill.ca/~jpineau

Overview of AI


- Typically three components:



Example AI system (1997): Chess playing

IBM Deep Blue defeated world champion Garry Kasparov.

- Perception: advanced features of the board.
- Actions: choose a move.
- Reasoning: search and evaluation of possible board positions.



<http://www-03.ibm.com>

Example AI system (2008): Poker playing

University of Alberta's Polaris defeats some of the world's best online pros.

- Perception: features of the game.
- Actions: choose a move.
- Reasoning: search and evaluation of possible moves, machine learning.



Match Number	Player	Amount Won	Player	Amount Won	Difference	Result
Remote 1	Matt Hawrilenko	+\$199500	IJay Palansky	-\$174000	+\$25500	Humans Win
Remote 2	Nick Grudzien	-\$2000	Kyle Hendon	-\$118000	-\$120000	Polaris Wins
Live 1	Nick Grudzien	-\$42000	Kyle Hendon	+\$37000	-\$5000	Draw
Live 2	Rich McRoberts	+\$89500	Victor Acosta	-\$39500	+\$50000	Humans Win
Live 3	Mark Newhouse	+\$251500	IJay Palansky	-\$307500	-\$56000	Polaris Wins
Live 4	Matt Hawrilenko	-\$60500	IJay Palansky	-\$29000	-\$89500	Polaris Wins

<http://poker.cs.ualberta.ca/>

Example AI system (2011): Jeopardy!

DISCOVER
M A G A Z I N E

Health & Medicine | Mind & Brain | Technology | Space | Human Origins | Living World | Environment

Technology / Computers

Top 100 Stories of 2011 #3: A Supercomputer Wins Jeopardy!

When IBM's game-playing computer trounced two trivia experts, its victory was hailed as a landmark for intelligent machines. A Jeopardy! champ explains why the real winners were humans.

by **Leeandra Keany**
From the **January-February special issue**; published online for subscribers only on December 29, 2011

Digg | Stumble | Like? | Share This



IBM's Watson computer system, powered by IBM POWER7, competes against Jeopardy!'s two most successful and celebrated contestants -- Ken Jennings and Brad Rutter.

COMP-424: Artificial intelligence 5 *Joelle Pineau*

Game playing


- One of the oldest, most well-studied domains in AI! **Why?**

Game playing

- One of the oldest, most well-studied domains in AI! **Why?**
 - People like them! People are good at playing them!
 - Often viewed as an indicator of intelligence.
 - State spaces are very large and complicated.
 - Sometimes there is stochasticity and imperfect information.
 - Clear, clean description of the environment.
 - Easy performance indicator.

“Games are to AI as Grand Prix racing is to automobile design”.

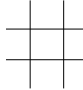
Start with an easy game: Tic-Tac-Toe



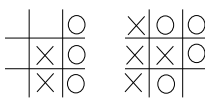
Defining a search problem for games

- **State space** S : all possible configurations of the domain.

- **Initial state** $s_0 \in S$: the start state

E.g. 

- **Goal states** $G \subset S$: the set of end states

E.g.  etc.

- **Actions** A : the set of moves available

Defining a search problem for games

- **Path**: a sequence of states and operators.

Defining a search problem for games

- **Path:** a sequence of states and operators.
- **Solution** of search problem: a path from s_0 to $s_g \in G$

Defining a search problem for games

- **Path:** a sequence of states and operators.
- **Solution** of search problem: a path from s_0 to $s_g \in G$
- **Utility:** a numerical value associated with a state (higher is better, lower is worse).
 - E.g. +1 if it's a win,
 - 1 if it's a loss,
 - 0 if it's a draw or game not terminated.

Representing search: Graphs and Trees

- Visualize the state space search in terms of a **graph**.

Representing search: Graphs and Trees

- Visualize the state space search in terms of a **graph**.
- Graph defined by a **set of vertices** and a **set of edges** connecting the vertices.
 - **Vertices** correspond to **states**.
 - **Edges** correspond to **actions**.

Representing search: Graphs and Trees


- Visualize the state space search in terms of a **graph**.
- Graph defined by a **set of vertices** and a **set of edges** connecting the vertices.
 - **Vertices** correspond to **states**.
 - **Edges** correspond to **actions**.
- We search for a solution by **building a search trees** and **traversing it to find a goal state**.

Search tree for Tic-Tac-Toe




We want to **find a strategy** (i.e. way of picking moves) that **wins the game**.

Search tree for Tic-Tac-Toe



We want to find a strategy (i.e. way of picking moves) that wins the game.

Search tree for Tic-Tac-Toe



Etc.

We want to find a strategy (i.e. way of picking moves) that wins the game.


Game search challenge

- Not quite the same as simple graph searching.
- There is an opponent! The opponent is malicious!
 - Opponent is trying to make things good for itself, and bad for us.
 - We have to simulate the opponent's decisions.

Game search challenge

- Not quite the same as simple graph searching.
- There is an opponent! The opponent is malicious!
 - Opponent is trying to make things good for itself, and bad for us.
 - We have to simulate the opponent's decisions.
- **Key idea:**
 - Define a **max player** (who wants to maximize the utility)
 - And a **min player** (who wants to minimize the utility.)

Example: Tic-Tac-Toe



COMP-102

21

Joelle Pineau

Minimax search

- Expand complete search tree, until terminal states have been reached and their utilities computed.

COMP-102


22

Joelle Pineau


Minimax search

- Expand complete search tree, until terminal states have been reached and their utilities computed.
- Go back up from leaves towards the current state of the game.
 - At each min node: backup the worst value among the children.
 - At each max node: backup the best value among the children.

A simple Minimax example



A simple Minimax example




COMP-102

25

Joelle Pineau

A simple Minimax example




COMP-102

26

Joelle Pineau

A simple Minimax example




COMP-102

27

Joelle Pineau

A simple Minimax example



COMP-102

28

Joelle Pineau

Properties of Minimax search

- Can we use minimax to solve any game?

Properties of Minimax search

- Can we use minimax to solve any game?
 - Solve Tic-Tac-Toe? Yes!
 - Solve chess? No.
- Why not?
 - Large number of actions possible (i.e. large branching factor) $b \approx 35$.
 - Path to goal may be very long (i.e. deep tree) $m \approx 100$
 - Large number of states!

Coping with resource limitations

- Suppose we have 100 seconds to make a move, and we can search 10^4 nodes per second.
 - Can only search 10^6 nodes!
(Or even fewer, if we spend time deciding which nodes to search.)
- Possible approach:
 - Only search to a pre-determined depth.
 - Use an evaluation function for the nodes where we cutoff the search.

Cutting the search effort

- Use an evaluation function to evaluate non-terminal nodes.
 - Helps us make a decision without searching until the end of the game.
- Minimax cutoff algorithm:
Same as standard Minimax, except stop at some maximum depth m and use the evaluation function on those nodes.

Evaluation functions

- An evaluation function $v(s)$ represents the “goodness” of a board state (e.g. chance of winning from that position).
 - Similar to a utility function, but approximate.
- If the features of the board can be evaluated **independently**, use a **linear combination**:

$$v(s) = f_1(s) + f_2(s) + \dots + f_n(s) \quad (\text{where } s \text{ is board state})$$
- This function can be given by the designer or learned from experience.

Example: Chess



Black to move
White slightly better



White to move
Black winning


- Evaluation function: $v(s) = f_1(s) + f_2(s)$

$$f_1(s) = w_1 * [(\# \text{ white queens}) - (\# \text{ black queens})]$$

$$f_2(s) = w_2 * [(\# \text{ white pawns}) - (\# \text{ black pawns})]$$

How precise should the evaluation fn be?

- Evaluation function is only approximate, and is usually better if we are close to the end of the game.
- Only the order of the numbers matter: payoffs in deterministic games act as an **ordinal utility function**.




Minimax cutoff in Chess

- How many moves ahead can we search in Chess?
 >> 10^6 nodes with $b=35$ allows us to search 4 moves ahead!
- Is this useful?


Minimax cutoff in Chess

- How many moves ahead can we search in Chess?
 - >> 10^6 nodes with $b=35$ allows us to search 4 moves ahead!
- Is this useful?
 - 4-moves ahead \approx novice player
 - 8-moves ahead \approx human master, typical PC
 - 12-moves ahead \approx Deep Blue, Kasparov
- Key idea:
 - Search few lines of play, but search them deeply. **Need pruning!**

α - β Pruning example



α - β Pruning example




COMP-102

39

Joelle Pineau

α - β Pruning example




COMP-102

40

Joelle Pineau

α-β Pruning example




COMP-102

41

Joelle Pineau

α-β Pruning example




COMP-102

42

Joelle Pineau

α - β Pruning example



α - β Pruning

- Basic idea: if a path looks worse than what we already have, ignore it.
 - If the best move at a node cannot change (regardless of what we would find by searching) then no need to search further!
- Algorithm is like Minimax, but keeps track of best leaf value for our player (α) and best one for the opponent (β)

Properties of α - β pruning

- Pruning does not affect the final result! You will not play worse than without it.
- Good move ordering is key to the effectiveness of pruning.
 - With **perfect ordering**, explore approximately $b^{m/2}$ nodes.
 - Means double the search depth, for same resources.
 - In chess: this is difference between novice and expert player.
 - With **bad move ordering**, explore approximately b^m nodes.
 - Means nothing was pruned.
 - Evaluation function can be used to order the nodes.

The α - β pruning demonstrates the value of reasoning about which computations are important!

Human or computer - who is better?

Checkers:

–

Othello:

–

Chess:

–

Backgammon:

–

Go:

–

Bridge:

–

Human or computer - who is better?

Checkers:

- 1994: Chinook (U.of A.) beat world champion Marion Tinsley, ending 40-yr reign.

Othello:

- 1997: Logistello (NEC research) beat the human world champion.
- Today: world champions refuse to play AI computer program (because it's too good).

Chess:

- 1997: Deep Blue (IBM) beat world champion Gary Kasparov

Backgammon:

- TD-Gammon (IBM) is world champion amongst humans and computers

Go:

- Human champions refuse to play top AI player (because it's too weak)

Bridge:

- Still out of reach for AI players because of coordination issue.
-

Jeopardy!

- In Winter 2011, Watson, a computer program created by IBM, made history by winning at Jeopardy!
 - Main innovation of Watson: ability to answer questions posed in natural language.
 - How it works:
 - Watson is much better at buzzing in than its human opponents.
 - Watson isn't connected to the internet, but had access to 4TB of stored information (incl. all of Wikipedia).
 - When given a question, it extracts keywords, looks in database for related facts, compiles list of answers, and ranks them by confidence.
 - See also:
 - <http://www.jeopardy.com/minisites/watson/>
 - <http://www.youtube.com/watch?v=12rNbGf2Wwo>
-

Take-home message

- Understand the basic components (state space, start state, end state, utility function, etc.) required to represent the types of games discussed today.
- Know how to build the search tree.
- Understand the *how* and *why* of Minimax, Alpha-beta pruning, and evaluation functions.
- Have some intuition for what makes certain games harder than others.